arXiv:1411.1594v1 [stat.ME] 6 Nov 2014

Joint estimation of K related regression models with simple L_1-norm penalties

E. Ollier[⋆] **and V. Viallon**[‡]

[⋆] *Université de Lyon, F-69622, Lyon, France*

[‡] *Université de Lyon, F-69622, Lyon, France; Université Lyon 1, UMRESTTE, F-69373 Lyon; IFSTTAR, UMRESTTE, F-69675 Bron.*

vivian.viallon@univ-lyon1.fr ed.ollier@gmail.com

Abstract

We propose a new approach, along with refinements, based on L_1 penalties and aimed at jointly estimating several related regression models. Its main interest is that it can be rewritten as a weighted lasso on a simple transformation of the original data set. In particular, it does not need new dedicated algorithms and is ready to implement under a variety of regression models, *e.g.*, using standard R packages. Moreover, asymptotic oracle properties are derived along with preliminary non-asymptotic results, suggesting good theoretical properties. Our approach is further compared with state-of-the-art competitors under various settings on synthetic data: these empirical results confirm that our approach performs at least similarly to its competitors. As a final illustration, an analysis of road safety data is provided.

1. Introduction

With the emergence of high-dimensional data, penalized regression models have now become standard, with such penalties as the L_q-norm or quasinorm of the parameter vector, for some $q \geq 0$. Underlying the use of these penalties, a very common assumption when working with moderate to high dimensional data is that the theoretical parameter vector $\boldsymbol{\beta}^*$ is *sparse*: only a small subset of its p components is expected to be non-null. Sparsity-inducing penalties, such as those relying on the L_1-norm, are especially useful in this context: the L_1-norm being convex it further leads to approaches that can generally be solved efficiently; they are often referred to as the lasso. Under appropriate conditions, lasso estimates especially attain optimal convergence rates (up to logarithmic terms): sparsity-inducing approaches are not only appealing for interpretation matters, but also because they can improve upon non-penalized procedures and lead to optimal estimates regarding estimation and/or prediction accuracy (see *e.g.*, Bickel et al. (2009); Bühlmann and van de Geer (2011); Dalalyan et al. (2014) for the linear regression case).

In many applications, the objective is actually to estimate several, K say, parameter vectors $\boldsymbol{\beta}_1^*, \ldots, \boldsymbol{\beta}_K^*$ corresponding to K related regression models; this problem is often referred to as multi-task learning in the literature (Argyriou et al., 2008; Maurer and Pontil, 2013; Negahban and Wainwright, 2011; Lounici et al., 2009; Jacob et al., 2009). Our motivating example corresponds to a standard situation in epidemiology or clinical research, where data come from several *strata* of the population (Gertheiss and Tutz, 2012; Viallon et al., 2014). Each stratum can correspond to a type or dosage of treatment, a given geographical area or can be defined by crossing age category and gender, etc. The relationship between the response variable and the covariates has then to be studied on each stratum. This relationship is rarely strictly the same over all the strata, but some homogeneity is generally expected. As a result, simple strategies consisting in estimating one model per stratum or one model on all the strata pooled together are generally sub-optimal in this context, and appropriate approaches should automatically adapt for the level of heterogeneity over the strata. Besides the sparsity assumption on each $\boldsymbol{\beta}_k^*$, vectors $\boldsymbol{\beta}_k^*$ are usually expected to be close to each other in some sense. The principle of multi-task learning strategies is then to account for the presumed similarity among the $\boldsymbol{\beta}_k^*$'s, while not masking potential heterogeneities, in order to improve estimation performance. Recent methods rely on penalties encouraging estimators of matrix $\mathbf{B}^* = (\boldsymbol{\beta}_1^*, \ldots, \boldsymbol{\beta}_K^*) \in \mathbb{R}^{p \times K}$ to exhibit some given structure: low-rank with the trace-norm penalty, group-structure with L_1/L_2 or L_1/L_∞ penalties, etc. To solve the corresponding optimization problem, dedicated algorithms are generally needed which limits their use and generalization by practitioners: most of the available algorithms enable to consider only the L_2 loss (for linear regression typically), and/or the logistic or hinge loss. In this paper we propose a new approach which is very intuitive and easy to implement under a variety of models, *i.e.*, for a variety of loss functions. Indeed, our approach reduces to a weighted lasso on a straightforward transformation of the data. As a result, any function or package enabling the resolution of the weighted lasso can directly be used for the implementation. In particular, the glmnet R package of Friedman et al. (2010) allows the treatment of linear, logistic, Poisson, multinomial and Cox models. The principle of our approach is very simple: we encourage sparsity within individual $\boldsymbol{\beta}_k^*$ and similarity between the $\boldsymbol{\beta}_k^*$'s. As will be shown below, we reach this objective by simply using L_1-norm penalties and our approach returns estimates for matrix $\mathbf{B}^*$ that are sums of a rank one matrix and a sparse one.

In the following Section 2, we formally introduce the setting we consider, describe our approach, state its connections with previous works, discuss its practical implementation and present theoretical results. Then, results from simulation studies are presented in Section 3, where comparisons are made with state-of-the-art methods. Finally, an illustration of our approach is provided on road safety data in Section 4, where $K = 18$ experts have been asked to determine drivers' responsibility in road traffic accidents and the objective is to assess expert agreement.

2. Methods

Although the main applications we have in mind concern data gathered from various strata of a population, our work falls into the more general multi-task learning context. For the sake of simplicity, we will use the terminology stratum everywhere, but everything can be extended by replacing stratum by task. In addition, methods as well as their theoretical properties will be presented in the linear regression model with no intercept for ease of notation. Extensions to generalized linear models (McCullagh and Nelder, 1989) are straightforward. How to deal with intercept terms in practice is briefly discussed in Section 2.2.2 below.

2.1 Notations

For any integer $m \geq 1$, we will denote by $[m]$ the set of values $\{1, \ldots, m\}$. Moreover, $\mathbf{0}_m$ and $\mathbf{1}_m$ will denote the vectors of size m with components all equal to 0 and 1 respectively, while $\mathbf{I}_m$ will stand for the $(m \times m)$ identity matrix. For any vector $\mathbf{x} = (x_1, \ldots, x_m) \in \mathbb{R}^m$, we further denote by $supp(\mathbf{x})$ its support (*i.e.*, $supp(\mathbf{x}) = \{j \in [m] : x_j \neq 0\}$) and we set $\|\mathbf{x}\|_q = (\sum_{j\in[m]} |x_j|^q)^{1/q}$ for any real number $q \in (0, \infty)$, $\|\mathbf{x}\|_\infty = \max_j |x_j|$ and $\|\mathbf{x}\|_0 = |supp(\mathbf{x})|$, where $|E|$ denotes the cardinality of the set E. For any set $E \subseteq [m]$, we will denote by $\mathbf{x}_E$ the vector of $\mathbb{R}^{|E|}$ with components $(x_j)_{j\in E}$.
Denote by $K \geq 1$ the number of strata and by n_k the number of observations in stratum $k \in [K]$, with $n = \sum_{k\in[K]} n_k$ the total number of observations. We will assume that for any $k \in [K]$ response vectors $\mathbf{y}^{(k)} = (y_1^{(k)}, \ldots, y_{n_k}^{(k)})^T \in \mathbb{R}^{n_k}$ are related to design matrices $\mathbf{X}^{(k)} = ({\mathbf{x}_1^{(k)}}^T, \ldots, {\mathbf{x}_{n_k}^{(k)}}^T)^T \in \mathbb{R}^{n_k \times p}$ according to the following linear model

$$\mathbf{y}^{(k)} = \mathbf{X}^{(k)}\boldsymbol{\beta}_k^* + \boldsymbol{\varepsilon}^{(k)}. \tag{1}$$

In this equation, vectors $\boldsymbol{\varepsilon}^{(k)} = (\varepsilon_1^{(k)}, \ldots, \varepsilon_{n_k}^{(k)})^T \in \mathbb{R}^{n_k}$ denote noise vectors: variables $\varepsilon_i^{(k)}$ will be assumed to be independent and identically distributed (i.i.d.) with $\mathbb{E}\varepsilon_i^{(k)} = 0$ and $\mathrm{Var}(\varepsilon_i^{(k)}) = \sigma^2$ for some unknown $\sigma^2 > 0$, for all $i \in [n_k]$ and $k \in [K]$. Finally, vectors $\boldsymbol{\beta}_k^* \in \mathbb{R}^p$ are the K parameter vectors of interest, to be estimated. In the context considered here, the $\boldsymbol{\beta}_k^*$'s are expected to be sparse and close to each other in some sense.

2.2 Proposed approach

2.2.1 Principle

Our proposal relies on the following decomposition for $\boldsymbol{\beta}_k^*$, $k \in [K]$:

$$\boldsymbol{\beta}_k^* = \overline{\boldsymbol{\beta}}^* + \boldsymbol{\gamma}_k^*. \tag{2}$$

Here $\overline{\boldsymbol{\beta}}^*$ describes what is "common" over the strata, while $\boldsymbol{\gamma}_k^*$ captures the variation in stratum k around $\overline{\boldsymbol{\beta}}^*$. Of course, there are infinitely many such decompositions, but

assuming that vectors $\boldsymbol{\beta}_k^*$ are close to each other, those minimizing the quantity $\sum_k \|\boldsymbol{\gamma}_k^*\|_q$, for some $q \geq 0$, are naturally appealing. They correspond to very natural choices for vector $\overline{\boldsymbol{\beta}}^* \in \arg\min_{\overline{\boldsymbol{\beta}}} \sum_k \|\boldsymbol{\beta}_k^* - \overline{\boldsymbol{\beta}}\|_q$. In particular, it is straightforward that choices $q = 0, 1$ and 2 correspond to decompositions where, for each $j \in [p]$, $\overline{\beta}_j^*$ is a mode, a median and the mean of $(\beta_{1,j}^*, \ldots, \beta_{K,j}^*)$, respectively.

The principle of our proposal is to obtain estimators of the form $\widehat{\boldsymbol{\beta}}_k = \widehat{\overline{\boldsymbol{\beta}}} + \widehat{\boldsymbol{\gamma}}_k$ with sparse $\widehat{\overline{\boldsymbol{\beta}}}$ and sparse $\widehat{\boldsymbol{\gamma}}_k$. For appropriate values $\lambda_1 \geq 0$ and $\lambda_{2,k} \geq 0$, our estimates are then defined as

$$(\widehat{\overline{\boldsymbol{\beta}}}, \widehat{\boldsymbol{\gamma}}_1 \ldots, \widehat{\boldsymbol{\gamma}}_K) \in \underset{\overline{\boldsymbol{\beta}}, \boldsymbol{\gamma}_1, \ldots, \boldsymbol{\gamma}_K}{\operatorname{argmin}} \left\{ \sum_{k \geq 1} \frac{\|\mathbf{y}^{(k)} - \mathbf{X}^{(k)}(\overline{\boldsymbol{\beta}} + \boldsymbol{\gamma}_k)\|_2^2}{2n} + \lambda_1 \|\overline{\boldsymbol{\beta}}\|_1 + \sum_{k \geq 1} \lambda_{2,k} \|\boldsymbol{\gamma}_k\|_1 \right\}. \tag{3}$$

Clearly, $\widehat{\overline{\boldsymbol{\beta}}} = \mathbf{0}$ and $\widehat{\boldsymbol{\gamma}}_k = \widehat{\boldsymbol{\beta}}_k$ for λ_1 large enough: the K problems are solved independently (the corresponding strategy will be referred to as IndepLasso in the sequel). On the other hand, we have $\widehat{\boldsymbol{\beta}}_k = \widehat{\overline{\boldsymbol{\beta}}}$ for all k if the $\lambda_{2,k}$'s are large enough: optimal solutions are those of the lasso run on all the data pooled together (the corresponding strategy will be referred to as IdentLasso in the sequel). This is a nice property since selecting appropriate values for λ_1 and the $\lambda_{2,k}$'s allows us to adapt to the extent of homogeneity over the strata. We refer to Section 3.1.1 below for a discussion on how to chose λ_1 and $\lambda_{2,k}$ in practice.
The optimization problem described in Equation (3) is of course equivalent to the minimization of

$$\sum_{k \geq 1} \frac{\|\mathbf{y}^{(k)} - \mathbf{X}^{(k)}\boldsymbol{\beta}_k\|_2^2}{2n} + \lambda_1 \{\|\overline{\boldsymbol{\beta}}\|_1 + \sum_{k \geq 1} \frac{\lambda_{2,k}}{\lambda_1} \|\boldsymbol{\beta}_k - \overline{\boldsymbol{\beta}}\|_1\} \tag{4}$$

over $\boldsymbol{\beta}_k \in \mathbb{R}^p$ for $k \in [K]$ and $\overline{\boldsymbol{\beta}} \in \mathbb{R}^p$. Because $\overline{\boldsymbol{\beta}}$ now only appears in the penalty terms, it is easy to see that, at optimum, $\widehat{\overline{\beta}}_j$ is a weighted and shrunk version of the median of $(\widehat{\beta}_{1,j}, \ldots, \widehat{\beta}_{K,j})$; we will denote it by $\text{WSmedian}_{\mu_{[K]}}(\widehat{\beta}_{1,j}, \ldots, \widehat{\beta}_{K,j})$ with $\mu_{[K]} = (\mu_1, \ldots, \mu_K)$ and $\mu_k = \lambda_{2,k}/\lambda_1$. If all μ_k's are equal and tend to 0, then $\text{WSmedian}_{\mu_{[K]}}$ tends to the constant function returning $\mathbf{0}_p$: this corresponds to the decomposition $\widehat{\boldsymbol{\beta}}_k = \widehat{\boldsymbol{\gamma}}_k$, and then to the IndepLasso strategy. On the other hand, if all μ_k's are equal but tend to infinity, then $\text{WSmedian}_{\mu_{[K]}}$ tends to the standard median function: this corresponds to the decomposition $\widehat{\boldsymbol{\beta}}_k = \widehat{\overline{\boldsymbol{\beta}}} + \widehat{\boldsymbol{\gamma}}_k$, where $\widehat{\overline{\beta}}_j = \text{median}(\widehat{\beta}_{1,j}, \ldots, \widehat{\beta}_{K,j})$. If the μ_k's are all equal to $1/\tau$, with $\tau \in \mathbb{N}$, then $\text{WSmedian}_{\mu_{[K]}}$ is a shrunk median in the sense that

$$\widehat{\overline{\beta}}_j = \text{WSmedian}_{1/\tau, \ldots, 1/\tau}(\widehat{\beta}_{1,j}, \ldots, \widehat{\beta}_{K,j}) = \text{median}(\mathbf{0}_\tau^T, \widehat{\beta}_{1,j}, \ldots, \widehat{\beta}_{K,j}).$$

A last interesting example is when for some $\ell \in [K]$ $\mu_\ell \to \infty$ and the other μ_k's are fixed (finite), for all $k \neq \ell$. Then $\text{WSmedian}_{\mu_{[K]}}(\widehat{\beta}_{1,j}, \ldots, \widehat{\beta}_{K,j}) \to \widehat{\beta}_{\ell,j}$, leading to the

decomposition $\widehat{\boldsymbol{\beta}}_k = \widehat{\boldsymbol{\beta}}_\ell + \widehat{\boldsymbol{\gamma}}_k$, with $\widehat{\boldsymbol{\gamma}}_\ell = \mathbf{0}_p$: this corresponds to considering stratum ℓ as the reference one (see Section 2.3.1 below). To sum up, each particular choice for the $\lambda_{2,k}/\lambda_1$ ratios "identifies" a particular common effect vector $\widehat{\overline{\boldsymbol{\beta}}}$ with $\widehat{\overline{\beta}}_j$ defined as $\text{WSmedian}_{\mu_{[K]}}(\widehat{\beta}_{1,j}, \ldots, \widehat{\beta}_{K,j})$, and our approach encourages solutions $(\widehat{\boldsymbol{\beta}}_1, \ldots, \widehat{\boldsymbol{\beta}}_K)$ with typically sparse vector of common effects $\widehat{\overline{\boldsymbol{\beta}}}$ and sparse vectors of differences $\widehat{\boldsymbol{\beta}}_k - \widehat{\overline{\boldsymbol{\beta}}}$.

2.2.2 Implementation: rewriting as a weighted lasso

A very attractive property of our approach is that it can be rewritten as a simple weighted lasso. To state this, first set $\mathscr{Y} = ({\mathbf{y}^{(1)}}^T, \ldots, {\mathbf{y}^{(k)}}^T)^T \in \mathbb{R}^n$ and introduce the quantities

$$\boldsymbol{\mathcal{X}} = \begin{pmatrix} \mathbf{X}^{(1)} & \mathbf{X}^{(1)} & \ldots & \mathbf{0} \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{X}^{(K)} & \mathbf{0} & \ldots & \mathbf{X}^{(K)} \end{pmatrix} \quad \text{and} \quad \boldsymbol{\theta} = \begin{pmatrix} \overline{\boldsymbol{\beta}} \\ \boldsymbol{\gamma}^{(1)} \\ \vdots \\ \boldsymbol{\gamma}^{(K)} \end{pmatrix},$$

which are elements of $\mathbb{R}^{n\times(K+1)p}$ and $\mathbb{R}^{(K+1)p}$ respectively. Further introduce the vector of weights $\boldsymbol{\omega} = (\mathbf{1}_p^T, (\lambda_{2,1}/\lambda_1)\mathbf{1}_p^T, \ldots, (\lambda_{2,K}/\lambda_1)\mathbf{1}_p^T)^T \in \mathbb{R}^{(K+1)p}$. Then, setting $\|\boldsymbol{\theta}\|_{1,\boldsymbol{\omega}} = \sum_{j=1}^{(K+1)p} \omega_j|\theta_j|$, the criterion to be minimized in Equation (3) can be rewritten as

$$\frac{\|\mathscr{Y} - \boldsymbol{\mathcal{X}}\boldsymbol{\theta}\|_2^2}{2n} + \lambda_1\|\boldsymbol{\theta}\|_{1,\boldsymbol{\omega}} \tag{5}$$

which is the criterion minimized in the weighted lasso. Of course, it is easy to show that this remains true under generalized linear models, Cox models, etc. As mentioned above, this is particularly interesting since it means that our approach can be implemented under a variety of models using available packages for the lasso, *e.g.*, the glmnet R package of Friedman et al. (2010).

Throughout the paper, only linear regression models with no intercept are considered for ease of notation. In practice however, intercept has generally to be included in the model. A first option consists in using the glmnet function with "intercept=FALSE" and replace $\mathbf{X}^{(k)}$ by $\mathbf{Z}^{(k)} = (\mathbf{1}_{n_k}, \mathbf{X}^{(k)}) \in \mathbb{R}^{n_k\times(p+1)}$ in the definition of $\boldsymbol{\mathcal{X}}$ above: this way, the absolute value of the common intercept term is penalized, and so are variations around this common intercept. In many situations, it makes more sense not to penalize the common intercept: an alternative option is then to use the glmnet function with "intercept=TRUE" and replace $\mathbf{X}^{(k)}$ by $\mathbf{Z}^{(k)}$ in the definition of $\boldsymbol{\mathcal{X}}$ everywhere except in the first p columns of $\boldsymbol{\mathcal{X}}$: this way, only variations around the common intercept (which corresponds to the median of the intercepts over the strata in this case) are penalized.

2.2.3 Related work

Decomposition (2) was first suggested in Evgeniou and Pontil (2004) who use the SVM machinery with L_2-norm penalties. For the application we have in mind, using sparsity-

inducing norms as the L_1-norm is more appealing. Indeed, detecting the subset of covariates that have a differential effect over the strata is often of primary interest (see our application in Section 4 below). By using the L_1-norm of the $\boldsymbol{\gamma}_k$'s our approach enjoys good properties regarding support recovery for both the $\boldsymbol{\beta}_k$'s and $\boldsymbol{\gamma}_k$'s (as will be shown below), which is of course not the case for methods based on L_2-norm penalties. Another advantage when using L_1-norms is that the common effect estimation is more robust (median versus mean).

Jalali et al. (2013) consider a closely related decomposition, this time for the parameter matrix $\mathbf{B} = (\boldsymbol{\beta}_1, \ldots, \boldsymbol{\beta}_K) \in \mathbb{R}^{p \times K}$. They write $\mathbf{B} = \mathbf{R} + \mathbf{S}$, where $\mathbf{R}$ and $\mathbf{S}$ are two $(p \times K)$ matrices with element (j,k) denoted by $r_j^{(k)}$ and $s_j^{(k)}$, respectively. Further denote by $\mathbf{r}^{(k)}$ and $\mathbf{s}^{(k)}$ the k-th column of matrices $\mathbf{R}$ and $\mathbf{S}$ respectively, so that $\boldsymbol{\beta}^{(k)} = \mathbf{s}^{(k)} + \mathbf{r}^{(k)}$, and by $\mathbf{r}_j$ and $\mathbf{s}_j$ the j-th row of matrices $\mathbf{R}$ and $\mathbf{S}$. Further set, for any matrix $\mathbf{M} \in \mathbb{R}^{p \times K} = (\mathbf{m}_1, \ldots, \mathbf{m}_p)^T$ and $q \geq 1$, $\|\mathbf{M}\|_{1,q} = \sum_{j \in [p]} \|\mathbf{m}_j\|_q$. Then their approach returns an estimate for $\mathbf{B}^*$ derived from minimizers of the following criterion:

$$\sum_k \frac{\|\mathbf{y}^{(k)} - \mathbf{X}^{(k)}(\mathbf{r}^{(k)} + \mathbf{s}^{(k)})\|_2^2}{2n} + \lambda_s \|\mathbf{S}\|_{1,1} + \lambda_r \|\mathbf{R}\|_{1,\infty}. \tag{6}$$

The term $\|\mathbf{S}\|_{1,1}$ encourages matrix $\mathbf{S}$ to be sparse (few elements are non-zero) while $\|\mathbf{R}\|_{1,\infty}$ encourages matrix $\mathbf{R}$ to have a row-wise group structure (few lines have zero entries): together, they encourage solutions $\widehat{\mathbf{B}}$ that can be written as the sum of a row-sparse matrix and a sparse one. As mentioned in Jalali et al. (2013), setting $d = \lfloor \lambda_r / \lambda_s \rfloor$, the combination of the two penalties leads to solutions $\widehat{\mathbf{R}}$ such that for any j with $\|\widehat{\mathbf{r}}_j\|_\infty > 0$ we have $|M_j| \geq d + 1$ where $M_j = \{k : |\widehat{r}_j^{(k)}| = \|\widehat{\mathbf{r}}_j\|_\infty\}$ (see their Lemma 2): in other words, rows that are not uniformly null have at least $d + 1$ components that have equal absolute values. In contrast, our approach returns estimates for $\mathbf{B}^*$ that is a sum of a rank one matrix (each row has p equal components instead of at least $d+1$ components of equal absolute values) plus a sparse matrix: it is therefore less flexible, but easier to interpret and, above all, to implement. From a theoretical point of view, Jalali et al. especially considered the case where $K = 2$ and covariates are generated from a standard multivariate Gaussian distribution. Further assuming that $n_1 = n_2$, they show that the number of samples needed to ensure support recovery with high probability (*i.e.*, the sample complexity) is inferior for Dirty, compared to both IndepLasso (where K lasso are run independently on the K strata) and the group-lasso strategy relying on the L_1/L_∞ penalty (Negahban and Wainwright, 2011). As shown Section 6.1.1 of the Appendix, the sample complexity of our approach is never superior to that of Dirty if in addition to the assumptions considered in Jalali et al., $\beta_{1,j}^* \beta_{2,j}^* \geq 0$ for all $j \in [p]$. Our empirical results presented in Section 3 further suggest that it may still be the case for $K > 2$. Stating this theoretically is out of the scope of the present paper and will be considered elsewhere.

2.3 Adaptive version and other refinement

2.3.1 Adaptive version

An adaptive version of our approach can be derived by selecting appropriate weights and then replacing the L_1 norms by weighted L_1-norms in Equation (3). Following the ideas of the adaptive lasso (Zou, 2006), these weights can be constructed from initial estimators of $\overline{\beta}_j$ and $\gamma_{k,j} = \beta_{k,j} - \overline{\beta}_j$, for $j \in [p]$ and $k \in [K]$. Denoting by $\widetilde{\boldsymbol{\beta}}_k$ initial estimates of $\boldsymbol{\beta}_k$ (*e.g.*, OLS estimates or MLEs if $n_k \gg p$ for all $k \in [K]$), and defining for any j,

$$\ell_j = \min\{k : \widetilde{\beta}_{k,j} \in \text{median}(\widetilde{\beta}_{1,j}, \ldots, \widetilde{\beta}_{K,j})\},$$

initial estimators for $\overline{\beta}_j$ and $\gamma_{k,j}$ are set to $\widetilde{\beta}_{\ell_j,j}$ and $\widetilde{\beta}_{k,j} - \widetilde{\beta}_{\ell_j,j}$. Note that because values $\widetilde{\beta}_{1,j}, \ldots, \widetilde{\beta}_{K,j}$ are generally all distinct, their median can be a range of values if K is even. In this case, we want to set the initial estimator of the common effect of covariate j to one of the values in $(\widetilde{\beta}_{1,j}, \ldots, \widetilde{\beta}_{K,j})$, hence the use of the minimum in the definition of ℓ_j (alternatively, the maximum could be used: there is no reason for favoring one or the other, at least a priori). Then setting, for all $k \in [K]$ and $j \in [p]$, $w_{k,j} = 1/|\widetilde{\beta}_{k,j}|^\rho$ and $\nu_{k,j} = 1/|\widetilde{\beta}_{k,j} - \widetilde{\beta}_{\ell_j,j}|^\rho$, for some $\rho > 0$ (a typical value is $\rho = 1$), the adaptive version of our approach consists in minimizing the following objective function

$$\sum_{k\geq 1} \frac{\|\mathbf{y}^{(k)} - \mathbf{X}^{(k)}(\overline{\boldsymbol{\beta}} + \boldsymbol{\gamma}_k)\|_2^2}{2n} + \lambda_1 \sum_{j=1}^{p} w_{\ell_j,j}|\overline{\beta}_j| + \sum_{k\geq 1} \lambda_{2,k} \sum_{j=1}^{p} \nu_{k,j}|\gamma_{k,j}| \Big\}.$$

But because $\nu_{\ell_j,j} = +\infty$ for all $j \in [p]$, this is equivalent to minimizing the criterion:

$$\sum_{k\geq 1} \frac{\|\mathbf{y}^{(k)} - \mathbf{X}^{(k)}\boldsymbol{\beta}_k\|_2^2}{2n} + \lambda_1 \sum_{j=1}^{p} w_{\ell_j,j}|\beta_{\ell_j,j}| + \sum_{k\geq 1} \lambda_{2,k} \sum_{j=1}^{p} \nu_{k,j}|\beta_{k,j} - \beta_{\ell_j,j}| \Big\}. \tag{7}$$

In other words, the adaptive version of our approach can be seen as a refinement of the following strategy, that is very common in epidemiology and clinical research when data come from several strata. Many practitioners would first select a reference stratum, say $\ell \in [K]$, and then use the decomposition: $\boldsymbol{\beta}_k = \boldsymbol{\beta}_\ell + \boldsymbol{\delta}_k$, for any $k \neq \ell$. Then, a lasso can be used and parameter estimates for each stratum can be obtained from minimizers of the following objective function:

$$\frac{1}{2n}\Big\{ \|\mathbf{y}^{(\ell)} - \mathbf{X}^{(\ell)}\boldsymbol{\beta}_\ell\|_2^2 + \sum_{k\neq\ell} \|\mathbf{y}^{(k)} - \mathbf{X}^{(k)}(\boldsymbol{\beta}_\ell + \boldsymbol{\delta}^{(k)})\|_2^2 \Big\} + \lambda_1\|\boldsymbol{\beta}_\ell\| + \sum_{k\neq\ell} \lambda_{2,k}\|\boldsymbol{\delta}^{(k)}\|_1. \tag{8}$$

This approach reduces to a standard lasso where to the original vector of covariates is augmented by interaction terms between the covariates and indicator functions describing membership to each stratum $k \neq \ell$. The adaptive version of our approach

enjoys two advantages compared to this strategy, referred to as InterLasso hereafter. First, the selection of the reference stratum is automatic with our approach and based on initial estimates of the $\boldsymbol{\beta}_k^*$,'s. Second, and above all, the reference stratum ℓ_j is covariate-specific. Whenever the initial estimates are consistent, ℓ_j belongs to the set $\mathcal{L}_j^* = \{k : \beta_{k,j}^* \in \text{median}(\beta_{1,j}^*, \ldots, \beta_{K,j}^*)\}$ with high probability for n large enough (if the n_k's all tend to ∞ as $n \to \infty$), and is therefore an appealing reference stratum for covariate j. As a result, our approach will generally yield models with lower complexity and hence better performance than the simple InterLasso (even if InterLasso can of course lead to better models than our approach in some situations; *e.g.*, if there exists a stratum ℓ for which $\beta_{\ell,j}^* \in \text{mode}(\beta_{1,j}^*, \ldots, \beta_{K,j}^*)$ for all $j \in [p]$, and this stratum ℓ is chosen as the reference one, and $\text{mode}(\beta_{1,j}^*, \ldots, \beta_{K,j}^*) \cap \text{median}(\beta_{1,j}^*, \ldots, \beta_{K,j}^*) = \emptyset$ for some $j \in [p]$).

2.3.2 Refinement

With our first approach, sparsity of vectors $\widehat{\boldsymbol{\beta}}_k = \widehat{\overline{\boldsymbol{\beta}}} + \widehat{\boldsymbol{\gamma}}_k$ is not directly encouraged, and is only "induced" by the sparsity of the common effects $\widehat{\overline{\boldsymbol{\beta}}}$ and of the variations $\widehat{\boldsymbol{\gamma}}_k$. For instance, if for some $a \neq 0$ and $j \in [p]$, $\beta_{\ell_j,j}^* = a$ and $\beta_{k,j}^* = 0$ for some $k \neq \ell$, our first approach will typically not return $\widehat{\beta}_{k,j} = 0$ (unless it also returns $\widehat{\beta}_{\ell_j,j} = 0$). We therefore propose a refined version which directly penalizes the L_1-norms of the $\boldsymbol{\beta}_k$'s. More precisely, a second set of estimators can be defined as minimizers of the following criterion

$$\sum_{k \geq 1} \Big\{ \frac{\|\mathbf{y}^{(k)} - \mathbf{X}^{(k)} \boldsymbol{\beta}_k\|_2^2}{2n} + \lambda_{1,k} \|\boldsymbol{\beta}_k\|_1 + \lambda_{2,k} \|\boldsymbol{\beta}_k - \overline{\boldsymbol{\beta}}\|_1 \Big\} \tag{9}$$

over $\boldsymbol{\beta}_k \in \mathbb{R}^p$ for $k \in [K]$ and $\overline{\boldsymbol{\beta}} \in \mathbb{R}^p$. In the sequel, this refined version will be referred to as M_2, while M_1 will refer to our first approach. Because $\overline{\boldsymbol{\beta}}$ now only appears in the last term of the objective function, it is clear that the j-th component of any optimal solution $\widehat{\overline{\boldsymbol{\beta}}}$ is such that $\widehat{\overline{\beta}}_j = \text{median}(\widehat{\beta}_{1,j}, \ldots, \widehat{\beta}_{K,j})$: for each covariate $j \in [p]$, its estimated common effect $\widehat{\overline{\beta}}_j$ corresponds to the median of its estimated effects over all the strata, irrespective to $\lambda_{1,k}/\lambda_{2,k}$ ratios.

The implementation of M_2 is however less trivial than that of our first approach. In Section 6.2 of the Appendix, we derive the dual formulation of the optimization problem described in Equation (9), which is shown to reduce to a standard optimization problem for the linear regression (quadratic programming) and the logistic regression (entropy maximization problem) and can therefore be solved with available optimization toolboxes; the hinge loss can also be easily treated.

An adaptive version of M_2 can further be recast in the generalized fused lasso framework (Höfling et al., 2010; Viallon et al., 2014). This makes its implementation straightforward with packages dedicated to the generalized fused lasso (*e.g.*, the FusedLasso R package of Höfling et al. (2010)). Recall the notations introduced in Section 2.3.1 above. The adaptive

version of M_2 consists in minimizing the following objective function

$$\sum_{k\geq 1}\Big\{\frac{\|\mathbf{y}^{(k)}-\mathbf{X}^{(k)}\boldsymbol{\beta}_k\|_2^2}{2n}+\lambda_{1,k}\sum_{j=1}^{p}w_{k,j}|\beta_{k,j}|+\lambda_{2,k}\sum_{j=1}^{p}\nu_{k,j}|\beta_{k,j}-\beta_{\ell_j,j}|\Big\}. \qquad (10)$$

Here again, we have $\nu_{\ell_j,j}=+\infty$ for all $j\in[p]$. To see the connection with the generalized fused lasso set $\mathscr{Y}=({\mathbf{y}^{(1)}}^T,\ldots,{\mathbf{y}^{(k)}}^T)^T\in\mathbb{R}^n$ as above, and introduce $\boldsymbol{\mathcal{X}}_F$ the $(n\times Kp)$ block diagonal matrix with k-th block equal to $\mathbf{X}^{(k)}$. Further define $\mathbf{b}=(\boldsymbol{\beta}_1^T,\ldots,\boldsymbol{\beta}_K^T)^T=(b_1,\ldots,b_{Kp})\in\mathbb{R}^{pK}$. Then, the criterion of Equation (10) can be rewritten as $\|\mathscr{Y}-\boldsymbol{\mathcal{X}}_F\mathbf{b}\|_2^2/(2n)+\lambda_1\|\mathbf{b}\|_1+\lambda_2\sum_{j_1\sim j_2}|b_{j_1}-b_{j_2}|$. Condition $j_1\sim j_2$ indicates that components b_{j_1} and b_{j_2} are connected in a particular graph that describes absolute differences that are penalized. Here, this graph consists of p star-graphs where the j-th star-graph has coefficient $\beta_{\ell_j,j}$ at its center, that is connected to each $\beta_{k,j}$ for $k\neq\ell_j$. The adaptive version of M_1 can also be seen as a particular case of the generalized fused lasso, using the same graph made of p star-graphs (our two adaptive versions penalize the same absolute differences), but where only the absolute value of the central node $|\beta_{\ell_j,j}|$ of each star-graph is penalized (while all the $|\beta_{k,j}|$'s are penalized in Equation (10) above). Finally note that the InterLasso strategy can also be seen as a special case of the generalized fused lasso: the underlying graph is made of p star-graphs with $\beta_{\ell,j}$ as the central node (instead of $\beta_{\ell_j,j}$), and only the $|\beta_{\ell,j}|$ terms are penalized, for $j\in[p]$.

These observations formally establish a connection with another strategy that was considered in Gertheiss and Tutz (2012) and Viallon et al. (2014). It will be referred to as CliqueFused, and it corresponds to a generalized fused lasso using this time a graph made of p cliques (instead of p star-graphs): for each covariate j, all absolute differences $|\beta_{k,j}-\beta_{k',j}|$ are penalized for all $k\neq k'$.

The rewriting of the adaptive version of M_2 as a special case of the generalized fused lasso shows that it is easy to implement with available packages, like the FusedLasso R packages for linear and logistic regression models; the same naturally holds for the adaptive version of M_1, but since it can still be written as a particular case of the (adaptive) lasso its implementation is faster using the glmnet function for instance. In addition asymptotic oracle properties for the adaptive version of M_2 (and M_1) readily follow as direct consequences of the results obtained in Viallon et al. (2014) (see Section 2.4 below); for the adaptive version of M_1, results obtained by Zou (2006) for the adaptive lasso can also be used.

2.4 Theoretical results

For the sake of brevity, we only present here a summary of asymptotic oracle properties for the adaptive versions of our approach. We refer to Section 6.1 in the Appendix for more details along with preliminary non-asymptotic results for our first approach M_1.

We consider the situation where p and K are held fixed (they do not increase with n). We further assume that, for each $k \in [K]$, matrix ${\mathbf{X}^{(k)}}^T\mathbf{X}^{(k)}/n_k$ converges to a positive definite matrix $\mathbf{C}^{(k)}$ as $n_k \to \infty$, and $n_k/n \to \kappa_k$ as $n \to \infty$, with $0 < \kappa_k < 1$ (i.e., *stratum* sizes all tend to infinity at the same rate). These assumptions essentially imply that MLEs $\widetilde{\beta}_{k,j}$ exist and are $\sqrt{n}$-consistent as $n \to \infty$; they will therefore be used in the definition of the weights in Equations (7) and (10). Our results essentially show that the adaptive version of M_2, say M_2^{ad}, enjoys *asymptotic oracle properties*. In particular, setting for any $j \in [p]$, $\ell_j^* = \min\{k : \beta_{k,j}^* \in \text{median}(\beta_{1,j}^*, \ldots, \beta_{K,j}^*)\}$ and $K_{\mathcal{A}^*,j}^* = \{k : \beta_{k,j}^* = \beta_{\ell_j^*,j}^*\}$ if $\beta_{\ell_j^*,j}^* \neq 0$, then all the parameters $\beta_{k,j}^*$ for $k \in K_{\mathcal{A}^*,j}^*$ are estimated by the common value $\widehat{\beta}_{\ell_j,j}^{(ad)}$, with probability tending to 1 as $n \to \infty$. This estimator has the same Gaussian limit distribution as that of the estimator we would obtain by pooling all data corresponding to covariate j and strata in $K_{\mathcal{A}^*,j}^*$ together. Moreover, M_2^{ad} is asymptotically optimal in situations where, for all $j \in [p]$, $\beta_{k_1,j}^* = \beta_{k_2,j}^*$ implies that either $\beta_{k_1,j}^* = \beta_{k_2,j}^* = 0$ or $\beta_{k_1,j}^* = \beta_{k_2,j}^* = \beta_{\ell_j^*,j}^*$, as in examples 1 to 3 of Figure 4 in the Appendix. Asymptotic oracle properties can easily be derived for method M_1^{ad} as well but, for M_1^{ad} to be optimal we must have in addition: for all $j \in [p]$, $\{\beta_{\ell_j^*,j}^* \neq 0\} \Rightarrow \{\forall k \in [K], \beta_{k,j}^* \neq 0\}$. For instance, in example 2 on Figure 4, the estimator of $\beta_{1,j}^*$ returned by M_1^{ad} can not be 0 (unless M_1^{ad} returns a zero estimate for $\beta_{k_j,j}^*$ as well). Therefore, M_1^{ad} would typically return a model with overall complexity higher than the theoretical one, hence be sub-optimal. On the other hand, contrary to the CliqueFused strategy, both our approaches use star-graphs and are therefore sub-optimal in situations where non-zero values in $(\beta_{1,j}^*, \ldots, \beta_{K,j}^*)$ consist of at least two groups of identical values (and possibly some distinct non-zero values). For instance, in example 4 of Figure 4, neither M_1^{ad} nor M_2^{ad} can return non-null equal values for components $\widehat{\beta}_{1,j}^{(ad)}$, $\widehat{\beta}_{2,j}^{(ad)}$, and $\widehat{\beta}_{3,j}^{(ad)}$ (nor for $\widehat{\beta}_{8,j}^{(ad)}$, $\widehat{\beta}_{9,j}^{(ad)}$, and $\widehat{\beta}_{10,j}^{(ad)}$), while CliqueFused typically would, for n large enough. However, CliqueFused may of course be outperformed by our approaches on finite samples. This was illustrated in the simulation study conducted in Viallon et al. (2014) in the single-task setting: they evaluated the robustness of the generalized fused lasso to graph misspecification and especially observed that the clique-based strategy, though asymptotically optimal, was outperformed by other graphs-based strategies on finite samples. In the present multi-task setting, if $\beta_{1,j}^*, \ldots, \beta_{K-1,j}^*$ are all equal and $\beta_{K,j}^*$ is different from the $K-1$ other ones, we would have $\ell_j \in [K-1]$ for n enough, and our approaches would only penalize $|\beta_{K,j} - \beta_{\ell_j,j}|$: they are therefore more likely to detect this difference than CliqueFused which penalizes all the differences $|\beta_{K,j} - \beta_{k,j}|$ for all $k \in [K-1]$.

3. Simulation Study

3.1 Competing Methods - Implementation

Our main objective here is to illustrate the two new approaches introduced in Section 2 and compare them with state-of-the-art competitors, in particular with those presented above which show some links with our proposal: CliqueFused and Dirty. For comparison, we further included IndepLasso and IdentLasso. The sparse group-lasso strategy relying on the $L_1 + L_1/L_\infty$ penalty and referred to as spGroupLasso hereafter is considered too (see Section 6.3 of the Appendix for a brief description of this approach and Negahban and Wainwright (2011) and Simon et al. (2013) for more details). Results obtained with the stepwise method described in Gertheiss and Tutz (2012) will also be presented in the low-dimensional example.

The glmnet R function of Friedman et al. (2010) was used to implement IndepLasso, IdentLasso as well as our approach M_1 (and its adaptive version). As for M_2, we used the Mosek Matlab toolbox (available at www.mosek.com) to solve the dual formulation presented in the Supplementaty Material. The FusedLasso R package of Höfling et al. (2010) was used to implement the CliqueFused approach and the adaptive version of M_2. The spGroupLasso strategy was implemented with the spams R package of Mairal et al. (2010), and we used the gvcm.cat R package for the stepwise approach of Gertheiss and Tutz (2012). Finally, we used the script of Jalali et al. (2013) (available at http://ali-jalali.com/index_files/L1Linf_LASSO.r) to implement Dirty. However, this script actually implements a revised version of Dirty. At each iteration ι of their algorithm, tuning parameters are divided by $\sqrt{\iota}$ in the coordinate gradient descent steps, making the soft-thresholding operator closer and closer to the hard-thresholding one as ι increases. This trick then returns solutions that are barely shrunk and is related in some way to the relaxed lasso of Meinshausen (2007), the adaptive lasso of Zou (2006) and L_q-penalization with $0 \leq q < 1$. For this reason, solutions returned by this script will be referred to as RevDirty. Note that the trick of RevDirty could of course be used for our approach if a lasso function existed with this trick; we do not know any such function or package though. RevDirty will therefore be included for the sake of completeness but the comparison with other strategies that do not use the corresponding trick is unfair. To implement the genuine Dirty as described in Jalali et al. (2013), we simply removed the "$/\sqrt{\iota}$" terms in the RevDirty script.

3.1.1 Practical selection of the tuning parameters

All methods presented above involve tuning parameters that need to be carefully selected in practice. Generally speaking, a predefined grid of λ values has first to be constructed. We refer to Section 6.4.3 of the Appendix for a complete description of the grid construction. Given grids of λ_1 and λ_2 values, cross-validation can be seen as the strategy of choice when p is large (*i.e.*, at least comparable to n) and/or only prediction accuracy matters. When $p \ll n$ and support recovery is of primary interest, a commonly preferred criterion is the

BIC computed with unbiased estimates: for the lasso, these unbiased estimates can be obtained by the OLS-Hybrid two-step strategy (Efron et al., 2004), which corresponds to a simplified version of the relaxed lasso with the particular choice $\phi = 0$ (Meinshausen, 2007). For the approaches considered in this paper, it can easily be extended. This criterion will be referred to as 2stepBIC in the sequel.

3.2 First simulation study

In the first simulation study, we consider the case where $K = 5$ and $n_k = 15p = 225$. For each stratum $k \in [K]$, each of the n_k rows (observations) of the design matrix $\mathbf{X}^{(k)}$ is generated under a multivariate Gaussian distribution $\mathcal{N}_p(\mathbf{0}_p, \mathbf{\Sigma})$, where the (j_1, j_2)-element of $\mathbf{\Sigma}$ is $2^{-|j_2-j_1|}$ (for $j_1, j_2 \in [p]$). For given values (SpGlob, SpSpec) $\in [0,1]^2$, each component $\overline{\beta}^*_j$ of vector $\overline{\boldsymbol{\beta}}^* \in \mathbb{R}^p$ is generated from a Bernoulli distribution with parameter SpGlob, while for each $k \in [K]$, each component of vector $\boldsymbol{\delta}^*_k \in \mathbb{R}^p$ is generated from a Bernoulli distribution with parameter SpSpec. The expected number of non-zero components in $\overline{\boldsymbol{\beta}}^*$ and $\boldsymbol{\delta}^*_k$ is then $p \times$ SpGlob and $p \times$ SpSpec respectively. Each vector $\boldsymbol{\beta}^*_k$ are then set to $\overline{\boldsymbol{\beta}}^* + \boldsymbol{\delta}^*_k$, with components either 0, 1 or 2. By making SpSpec vary, we make the level of heterogeneity vary over the strata, while SpGlob enables the control of the "common" effects sparsity level. Observe that components $\overline{\beta}^*_j$ do not necessarily correspond to median$(\beta^*_{1,j}, \ldots, \beta^*_{K,j})$ (they do in some cases, *e.g.*, when SpSpec= 0, but not always). Given $\boldsymbol{\beta}^*_k$ and $\mathbf{X}^{(k)}$, vector $\mathbf{y}^{(k)}$ is generated from a multivariate gaussian distribution $\mathcal{N}_{n_k}(\mathbf{X}^{(k)}\boldsymbol{\beta}^*_k, \sigma^2\mathbf{I}_{n_k})$. By varying the noise variance σ^2, we can make the *signal-to-noise-ratio* (SNR) vary. For each simulation design, corresponding to a given triple of values (SpGlob, SpSpec, SNR), 100 replications are performed, and results presented and discussed below correspond to averages over these 100 replications. Methods are evaluated according to their prediction accuracy (Figure 1), measured by $\log(\sum_{k\in[K]} \|\mathbf{X}^{(k)}(\boldsymbol{\beta}^*_k - \widehat{\boldsymbol{\beta}}_k)\|_2^2)$, and support recovery accuracy (Figure 2), that measures the ability to correctly recover zero and non-zero elements of matrix $\mathbf{B}^* = (\boldsymbol{\beta}^*_1, \ldots, \boldsymbol{\beta}^*_K)$. Tuning parameters are selected with the 2StepBIC described above.

In Figure 1, the first row corresponds to SpSpec= 0, so that all $\boldsymbol{\beta}^*_k$'s are equal. In this case, the optimal strategy is of course IdentLasso but M_2, spGroupLasso, M_1 and Dirty all perform as well (or nearly as well) as this optimal strategy. The stepwise approach and CliqueFused are less efficient, but still perform better than IndepLasso. On the other hand, when SpGlob is low and SpSpec is high (last row, first two columns), the K true models have not much in common and IndepLasso is nearly-optimal: in this case again, the multi-task learning strategies perform nearly as well as the optimal strategy. In addition, and as expected, prediction performance of all methods tend to decrease as SpGlob and/or SpSpec increases, that is as the true model complexity increases. Overall, M_1 and Dirty show very similar performance, irrespective to the true model complexity. They perform the best for SpSpec$\neq$ 0, closely followed by the other multi-task learning strategies. For SpGlob= 0.4 and SpSpec$\neq$ 0, we observed moderate performance for both M_2 and spGroupLasso.

As shown in Section 6.4.3 of the Appendix for method M_2, alternative choices for the initial grid of (λ_1, λ_2) values lead to better results, suggesting that our proposal for the construction of this grid might not always be optimal for M_2 (and spGroupLasso) and that more theoretical effort might be needed to get a better initial grid.

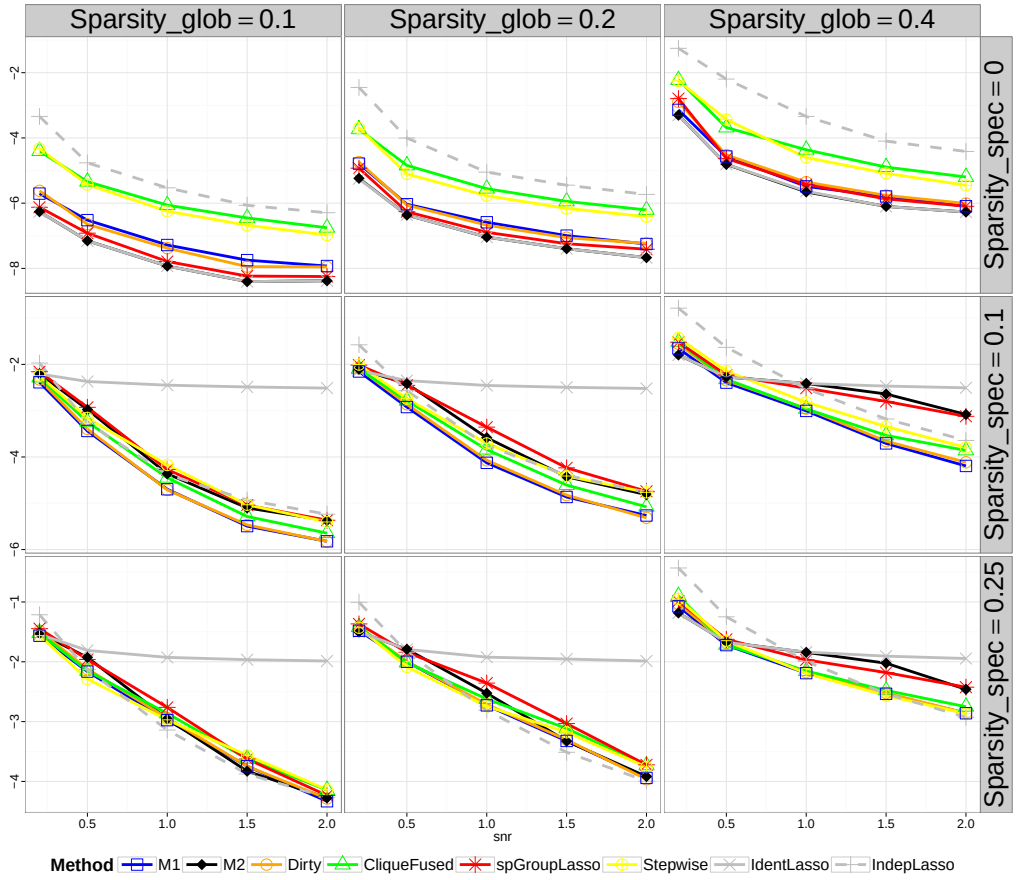

Figure 1: Results from the first simulation study: prediction accuracy.

Turning our attention to support recovery, Figure 2 presents results obtained with Dirty and RevDirty, M_1, and the adaptive versions of M_1, M_2 and CliqueFused. See Figure 6 in the Appendix for the corresponding analysis of the methods compared on Figure 1: overall, conclusions for these methods are very similar to those drawn from the analysis of the prediction performance. Weights for adaptive versions were derived from initial OLS estimates of the $\boldsymbol{\beta}_k^*$. The comparison of M_1 with its adaptive version $M_1^{(ad)}$ clearly illustrates the potential gain on support recovery accuracy when using adaptive weights (no such gain was observed on prediction accuracy; results not shown). Moreover, adaptive

versions of CliqueFused and M_2 attain nearly the same performance as $M_1^{(ad)}$. Overall, RevDirty performs the best, but (i) the comparison is unfair since RevDirty is more than an adaptive version of Dirty and (ii) adaptive versions of M_1 and M_2 performs only slightly worse in most cases.

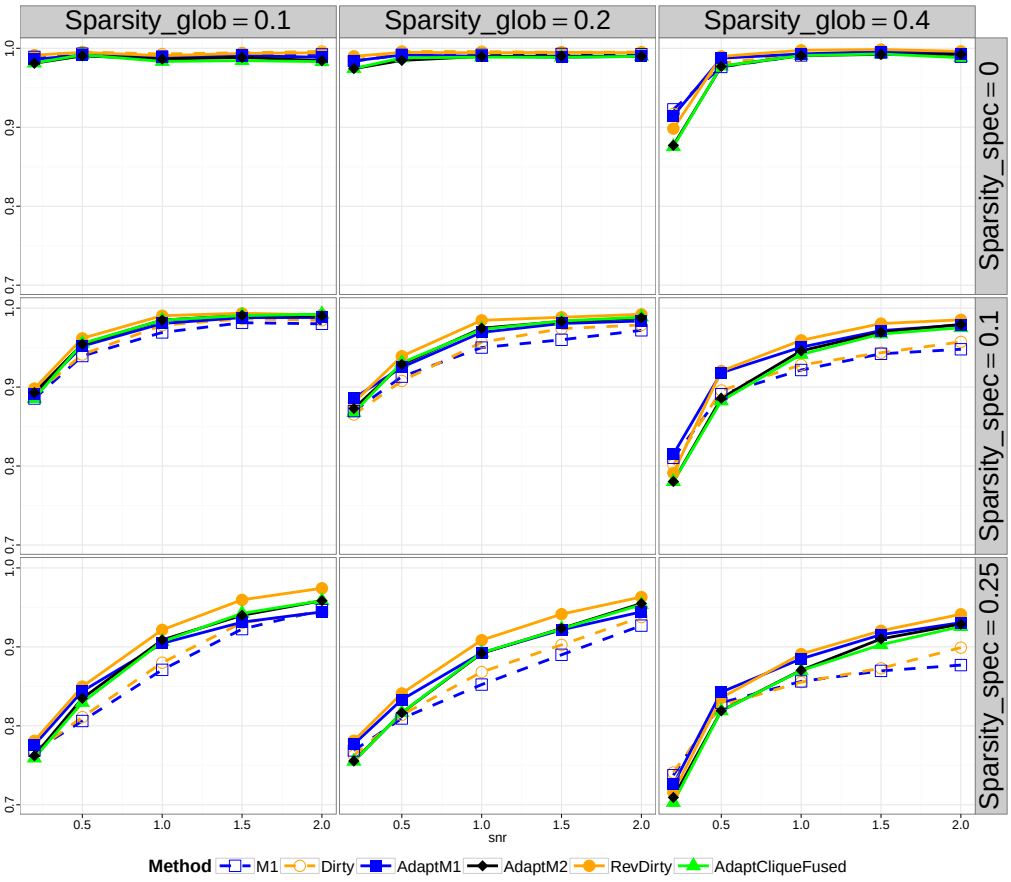

Figure 2: Results from the first simulation study: support recovery accuracy.

3.3 Second simulation study

The objective of the second simulation study was to describe the methods performance when the dimension p is of the order of n_k and selection of tuning parameters is done by cross-validation. In view of the results of the first simulation study, and to save computational times, IdentLasso, IndepLasso and CliqueFused were not included in this analysis. In addition, a relaxed version of M_1 was included, following the idea of the relaxed lasso

(Meinshausen, 2007) (we made the ϕ-value vary over the sequence $\{0, 0.1, \ldots, 0.9, 1\}$). We considered the case where $K = 5$ and $p = n_k = 100$. Vectors $\overline{\boldsymbol{\beta}}$, and $\boldsymbol{\delta}_k^*$, for $k \in [K]$ are constructed as follows. Only the first 10 components of these vectors can be non-zero. The sparsity of the common effect $\overline{\boldsymbol{\beta}}$ is fixed to 4, with non-zero components all equal to 1. Three levels of heterogeneity are considered: complete homogeneity (all the $\boldsymbol{\delta}_k^*$'s are equal to $\mathbf{0}_p$), low heterogeneity (the number of non-zero elements of matrix $\boldsymbol{\Delta}^* = (\boldsymbol{\delta}_1^*, \ldots, \boldsymbol{\delta}_K^*)$ is 5, each non-zero element being ± 1 with probability 1/2), and moderate heterogeneity (the overall sparsity of matrix $\boldsymbol{\Delta}^*$ is 15, each non-zero element being ± 1 with probability 1/2). Design matrices $\mathbf{X}^{(k)}$ and response vectors $\mathbf{y}^{(k)}$ are generated as in the first simulation study, except we now set $\boldsymbol{\Sigma} = \mathbf{I}_p$. As in the first simulation study, we make the SNR vary by varying the noise level.

Results are presented in Figure 3. Because Dirty and RevDirty take a very long time to run in this second simulation study, results correspond to averages over 20 replications only. Methods are evaluated according to their prediction accuracy, measured by $\log(\sum_{k\in[K]} ||\mathbf{X}^{(k)}(\boldsymbol{\beta}_k^* - \widehat{\boldsymbol{\beta}}_k)||_2^2)$, their estimation accuracy (measured by $\sum_{k\in[K]} \|\boldsymbol{\beta}_k^* - \widehat{\boldsymbol{\beta}}_k)\|_2^2)/(Kp)$), and their ability to recover the true support (measured in this highly sparse setting by the F1-score, defined as the harmonic mean of precision and recall). Because estimators selected by cross-validation may have many tiny non-zero components, hard-thresholding these estimators can improve their performance regarding support recovery. F1-score for these hard-thresholded versions are therefore also presented, with threshold set to 0.25 and 0.5. By first comparing crude versions (no adaptive weights, nor relaxation nor the Revised Dirty trick), M_1 and M_2 generally perform at least similarly to Dirty and spGroupLasso. The relaxed and adaptive versions of M_1 generally achieve better performance than the crude M_1, especially in cases of complete homogeneity or low heterogeneity. In these cases, these versions also outperform RevDirty in terms of prediction and estimation accuracy. As for the F1-score, RevDirty is clearly the best: only in the case of complete homogeneity the relaxed version of M_1 achieves similar performance. Note however that the comparison is again unfair, since RevDirty is more than an adaptive or a relaxed version of Dirty. After the hard-thresholding step, differences among methods in terms of F1-scores are much narrower: focusing on low values of the SNR, the different versions of our approaches appear very competitive. (Our results confirm that RevDirty returns nearly unshrunk estimates and then fewer tiny non-zero components than "purely" L_1-based approaches: it does fewer mistakes (its F1-score is higher), but apparently bigger mistakes since its prediction and estimation accuracies can be outperformed by the relaxed or adaptive versions of M_1.)

4. Application on road safety data

One major objective in road safety is to determine factors associated with being responsible of a road traffic accident. To do so, epidemiologists typically use data sets of car crashes, in which drivers' responsibility has been determined by experts. Of course, it is crucial that

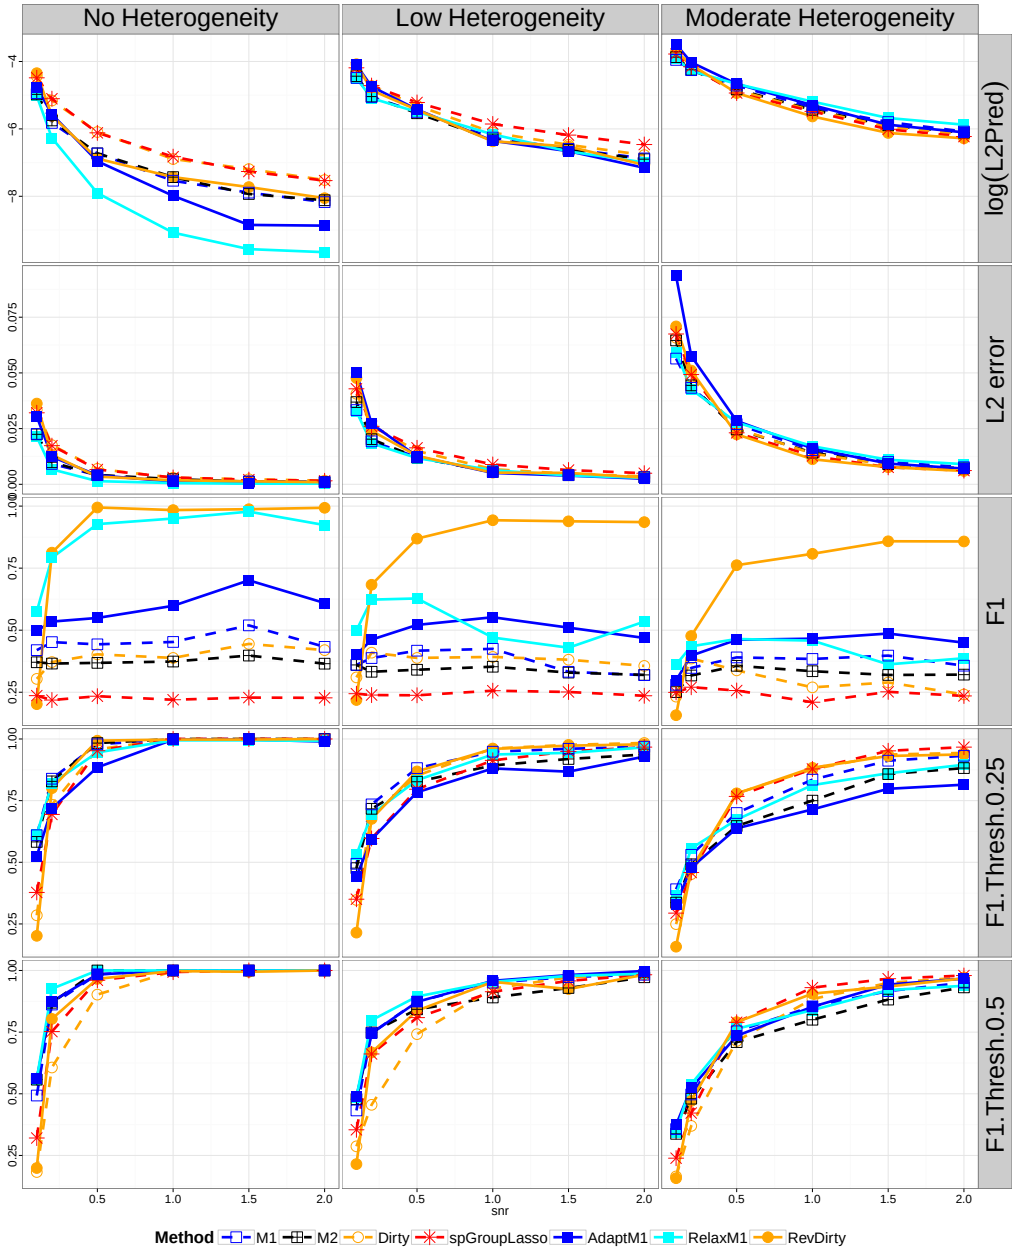

Figure 3: Results from the second simulation study with $n_k = p = 100$ and $K = 5$.

these experts use a proper and consistent "rule" to determine responsibility. Consistency especially implies that two experts should agree on drivers' responsibility. Level of agreement can be measured by Cohen's κ or other standard agreement criteria, but this would require each driver's responsibility to be rated by at least two experts. Here, we will use a joint estimation strategy to assess experts agreement in situations where each driver's responsibility is only rated once by one expert. As will be made clearer below, our strategy will also help us to check the relevance of the rule used by experts for the determination of drivers' responsibility.

Our dataset consists of $n = 5772$ road traffic accidents involving two vehicles (cars, buses, trucks, two-wheeled motor vehicles, bikes, ...). These data have been collected in the VOIESUR project (see www.agence-nationale-recherche.fr/projet-anr/?tx_lwmsuivibilan_pi2[CODE]=ANR-11-VPTT-0007). For each accident, the responsibility of the two drivers was rated by one of the eighteen recruited experts. To do so, experts had to analyze police reports established after the accident. Prior to this, experts underwent a period of training where basic recommandations were given: for instance, driving under the influence (of alcohol or any other drug) should not directly be taken into account by experts, but only "incorrect operations" (lane departure, ...) should. This is of course critical because, otherwise, evaluating the association between alcohol consumption and being responsible of a car crash with this dataset would lead to overestimation. Experts were asked to rate responsibility from 0 (not responsible at all) to 5 (totally responsible). For the present analysis, drivers with rate 0 or 1 were considered as non-responsible, drivers with rate 4 or 5 responsible, and drivers with rate 3 (unclear responsibility) were excluded. Besides drivers' responsibility, 205 binary covariates are at our disposal, describing drivers' characteristics (age, gender, profession, alcohol and other drugs consumptions, estimated speed, lane departure, ...), their antagonists' characteristics, meteorological conditions, etc... The rule used by each expert to determine drivers' responsibility can be described by the probability of being considered as responsible by this particular expert given the 205 covariates. Then, sparse logistic regression can be used for the estimation of this probability. If experts all use the same rule, then theoretical parameter vectors $\boldsymbol{\beta}_k^*$ of each logistic model ($k = 1, \ldots, 18$) should be equal. To check whether this is the case, we propose to jointly estimate the 18 corresponding sparse logistic models using the two approaches described in this paper (both the crude and refined versions). Each method returns a $(206, 18)$ matrix where each columns contains the parameter vector (including the intercept) of one stratum, *i.e.* one expert. We evaluated both methods on a predefined grid of $(\lambda_1, \lambda_{2,k})$ values and retained the model minimizing the 2stepBIC. Contrary to situations considered in our simulation study, the n_k's are not constant here and several strategies can be put forward for determining the grid. The most simple one consists in using $\lambda_{2,k} = \lambda_2$ and proceed as described in Section 3.1.1. As mentioned in Section 3.1.1 (and as suggested by the preliminary non-asymptotic analysis of Section 6.1.2 of the Appendix), another strategy for method M_1 consists in first standardizing columns of matrix $\boldsymbol{\mathcal{X}}$, then setting $\lambda_{2,k} = \lambda_2$ and finally proceeding again

as described in Section 3.1.1. Both strategies were applied for method M_1 and they led to similar results.

With either method, good concordance was observed between supports of vectors $\widehat{\boldsymbol{\beta}}_k$: only two [resp. three] elements of matrix $\widehat{\boldsymbol{\Delta}} = (\widehat{\boldsymbol{\delta}}_1, \ldots, \widehat{\boldsymbol{\delta}}_{18})$ were non-null when using M_1 [resp. M_2]. Moreover, non-zero elements returned by M_2 have small absolute values ($\leq$ 0.2). This suggests good agreement among the 18 experts and we can consider they use very similar rules (if not equal) to determine drivers' responsibility. An interesting result is that elements in $\widehat{\boldsymbol{\Delta}}$ related to driving under the influence were all zero: this suggests that all experts account for alcohol consumption in the same way.

Now, looking at the common vector $\widehat{\overline{\boldsymbol{\beta}}}$ can further help us to check the adequacy of the common rule used by the experts. We will pay a particular attention on the association between alcohol consumption and the probability of being considered responsible by experts. Indeed for road safety matters, and as recalled to experts during the training phase, the only fact of driving under the influence should not imply to be (nor increase the risk of being) considered as responsible. If components of vectors $\widehat{\overline{\boldsymbol{\beta}}}$ corresponding to alcohol consumption are non-null, then the adequacy of the rule may be questioned. Method M_1 returns a vector $\widehat{\overline{\boldsymbol{\beta}}}$ with 74 non-null components, while M_2 returned a slightly sparser one with 60 non-null components. Overall, variables that are the most associated to a high probability of being responsible are as expected: a driver who drove (or ran) away, and/or went through a stop sign or a red light, crossed the lane, etc. is likely to be considered as responsible of the accident. Regarding alcohol consumption, both M_1 and M_2 return two non-zero components in vector $\widehat{\overline{\boldsymbol{\beta}}}$: those corresponding to the highest level of alcohol (between 1.2 and 2 g/L, and above 2g/L respectively). Two interpretations can be put forward. A pessimistic one is that experts did not follow recommandations and usually considered these levels of alcohol as indicators of drivers' responsibility. Another more optimistic interpretation is that drivers presenting these levels of alcohol may have behaviors that can not be captured by a simple linear combination of the 205 available binary variables (either because interactions should be considered or, more likely, because some descriptors are missing). Choosing between these two interpretations is out of the scope of the present analysis and will be the focus of future work.

Results from our approaches were compared to a more standard strategy in epidemiology which consists in using the SAS LOGISTIC procedure along with STEPWISE option to select significant covariates (including interaction terms between the 205 original covariates and indicator functions representing each expert). The STEPWISE procedure was not able to run with the whole set of interaction terms. We then had to focus on the 205 original covariates plus interaction terms involving alcohol-related covariates. The results returned by SAS were consistant with those obtained with our approaches. More precisely, no interaction term was retained by the STEPWISE procedure confirming that all experts account for alcohol in the same way when determining drivers' responsibility. Moreover, the STEPWISE retained terms corresponding to the two highest levels of alcohol, as our

approaches did. This example illustrates situations where our approaches, though very simple to implement, extend the scope of standard procedures used in epidemiology.

5. Discussion

In this paper we present a new approach, along with refinements, aimed at jointly modeling several related regression models. Although most of the presentation was made in the situation where these models correspond to several strata of a population, our approach can be used in the more general multi-task learning setting. Also, for ease of notation mostly linear regression models were considered in this paper. As illustrated in the application of Section 4, extension to other models is however straightforward. This is particularly true for our first approach and its adaptive or relaxed version, since they all can be rewritten as simple weighted lasso problems: in particular, linear, logistic, Poisson and Cox models can be treated thanks to the glmnet R package (Friedman et al., 2010). Other functions exist for other models, like the clogitL1 R package for conditional logistic models of Reid and Tibshirani (2014), and allow for further extensions of our approach.

Our approach naturally connects with several strategies formerly proposed in the literature. First, it is a simplified version of the Dirty models proposed by Jalali et al. (2013) in the linear regression setting, and is much easier to extend to other regression models. On the designs considered in our simulation study, it still performs similarly to the Dirty models and we also exhibited situations where its theoretical sample complexity is at least as good as that of the Dirty models. Second, adaptive versions of our approaches are special cases of the generalized fused lasso. This connection is particularly appealing for interpretation matters since it allows for a simple comparison between CliqueFused, the InterLasso and our two strategies. First, CliqueFused uses cliques, while InterLasso and our two approaches use star-graphs. Second, for all $j \in [p]$, the center of the star-graphs are set to $\beta_{\ell,j}$ when using InterLasso with reference stratum ℓ, while when using adaptive versions of our approaches centers are adaptively set to $\beta_{\ell_j,j}$: the reference stratum is covariate-specific, automatically selected from the initial estimates, and it corresponds (with probability tending to 1) to one of the stratum in $K_j^* = \{k : \beta_{k,j}^* = \beta_{\ell_j^*,j}^*\}$ if the initial estimates are consistent. Finally, besides the structure of the graph (edges corresponding to differences $\beta_{k,j} - \beta_{k',j}$ that are penalized), these four strategies exhibit differences according to which nodes in the graph (coefficients $\beta_{k,j}$) are directly encouraged to be sparse: they are all penalized when using CliqueFused or M_2, while only the $\beta_{\ell,j}$ for InterLasso and the $\beta_{\ell_j,j}$ for M_1 are penalized, for $j \in [p]$.

From a theoretical point of view, the connection with the generalized fused lasso allowed us to derive asymptotic oracle properties for our adaptive versions. Preliminary non-asymptotic results were also stated, under strong but not necessary conditions. Regarding prediction performance, these preliminary results might be extended to more general settings by adapting the recent results of Dalalyan et al. (2014), who study prediction performance of the lasso in the presence of correlated designs. As for support recovery

of the $\boldsymbol{\beta}_k^*$'s, adapting the dual-witness idea of Wainwright (2009) is a promising lead (as Jalali et al. (2013) did for the Dirty models). Other extensions of this work could rest on greedy algorithms or recent MCMC-based approaches (such as the exponential screening of Rigollet and Tsybakov (2012)), for which optimal prediction bounds have been obtained with no particular assumption on the design matrix.

References

A. Argyriou, T. Evgeniou, and M. Pontil. Convex multi-task feature learning. *Machine Learning*, 73(3):243–272, 2008.

P. J. Bickel, Y. Ritov, and A. B. Tsybakov. Simultaneous analysis of lasso and dantzig selector. *The Annals of Statistics*, 37(4):1705–1732, 2009.

P. Bühlmann and S. van de Geer. *Statistics for High-Dimensional Data: Methods, Theory and Applications*. Berlin : Springer-Verlag, 2011.

A. S. Dalalyan, M. Hebiri, and J. Lederer. On the prediction performance of the lasso. *arXiv preprint arXiv:1402.1700*, 2014.

B. Efron, T. Hastie, I. Johnstone, and R. Tibshirani. Least angle regression. *The Annals of statistics*, 32(2):407–499, 2004.

T. Evgeniou and M. Pontil. Regularized multi–task learning. In *Proceedings of the tenth ACM SIGKDD International Conference on Knowledge Discovery and Data mining*, pages 109–117. ACM, 2004.

J. Friedman, T. Hastie, and R. Tibshirani. Regularization paths for generalized linear models via coordinate descent. *Journal of statistical software*, 33(1):1, 2010.

J. Gertheiss and G. Tutz. Regularization and model selection with categorial effect modifiers. *Statistica Sinica*, 22:957–982, 2012.

H. Höfling, H. Binder, and M. Schumacher. A coordinate-wise optimization algorithm for the Fused Lasso. *Arxiv preprint arXiv:1011.6409*, 2010.

L. Jacob, J.P. Vert, and F. Bach. Clustered multi-task learning: A convex formulation. In *Advances in neural information processing systems*, pages 745–752, 2009.

A. Jalali, P. Ravikumar, and S. Sanghavi. A dirty model for multiple sparse regression. *IEEE Transactions on Information Theory*, 59(12):7947–7968, 2013.

K. Lounici. Sup-norm convergence rate and sign concentration property of lasso and dantzig estimators. *Electronic Journal of statistics*, 2:90–102, 2008.

K. Lounici, M. Pontil, A.B. Tsybakov, and S. van de Geer. Taking advantage of sparsity in multi-task learning. In *COLT'09*, 2009.

J. Mairal, F. Bach, J. Ponce, and G. Sapiro. Online Learning for Matrix Factorization and Sparse Coding. *The Journal of Machine Learning Research*, 11:19–60, 2010.

A. Maurer and M. Pontil. Excess risk bounds for multitask learning with trace norm regularization. In *COLT'13*, pages 55–76, 2013.

P. McCullagh and J. A. Nelder. *Generalized Linear Models. 2nd ed.* New-York : Chapman & Hall, 1989.

N. Meinshausen. Relaxed lasso. *Computational Statistics and Data Analysis*, 52(1):374–393, 2007.

S. Negahban and M. Wainwright. Simultaneous support recovery in high dimensions: Benefits and perils of block-regularization. *IEEE Transactions on Information Theory*, 57(6):3841–3863, 2011.

S. Reid and R. Tibshirani. Regularization paths for conditional logistic regression: The clogitl1 package. *Journal of Statistical Software*, 58(12), 2014.

P. Rigollet and A. B. Tsybakov. Sparse estimation by exponential weighting. *Statistical Science*, 27(4):558–575, 2012.

N. Simon, J. Friedman, T. Hastie, and R. Tibshirani. A sparse-group lasso. *Journal of Computational and Graphical Statistics*, 22(2):231–245, 2013.

S. van de Geer and P. Bühlmann. On the conditions used to prove oracle results for the lasso. *Electronic Journal of Statistics*, 3:1360–1392, 2009.

V. Viallon, S. Lambert-Lacroix, H. Höfling, and F. Picard. On the robustness of the generalized fused lasso to prior specifications. *Statistics and Computing (To Appear)*, 2014.

M. Wainwright. Sharp thresholds for high-dimensional and noisy sparsity recovery using-constrained quadratic programming (lasso). *IEEE Transactions on Information Theory*, 55(5):2183–2202, 2009.

H. Zou. The Adaptive Lasso and Its Oracle Properties. *Journal of the American Statistical Association*, 101(476):1418–1429, 2006.

6. Appendix: Supplementary Materials

6.1 Theoretical results

In this section, we derive preliminary non-asymptotic properties for our first approach M_1 (Sections 6.1.1 and 6.1.2) as well as asymptotic oracle properties for our adaptive versions (especially for M_2, see Section 6.1.3).

6.1.1 Sample complexity under independent Gaussian designs, when $K = 2$ and $\beta^*_{1,j}\beta^*_{2,j} \geq 0$ for all $j \in [p]$.

Here we consider the case where $K = 2$. Working under the assumptions considered in Jalali et al. (2013), it is easy to show that the sample complexity of M_1 is at most the one of the Dirty models, under the additional assumption that $\beta^*_{1,j}\beta^*_{2,j} \geq 0$ for all $j \in [p]$. Indeed, for any given pair of potential estimates $\beta_{1,j}$ and $\beta_{2,j}$, the penalty term involved in the Dirty approach for covariate j can be written

$$\phi(r_1, r_2) = \lambda_1 \max(|r_1|, |r_2|) + \lambda_2\{|\beta_{1,j} - r_1| + |\beta_{2,j} - r_2|\},$$

while for M_1 it reduces to

$$\bar{\phi}(r) = \lambda_1|r| + \lambda_2\{|\beta_{1,j} - r| + |\beta_{2,j} - r|\} = \phi(r, r).$$

Moreover, for any positive λ_1, λ_2 and any $\beta_{1,j}, \beta_{2,j}$ it is easy to show that $\min_{r_1,r_2} \phi(r_1, r_2) = \min_r \bar{\phi}(r)$ if $\beta_{1,j}\beta_{2,j} \geq 0$ and $\min_{r_1,r_2} \phi(r_1, r_2) < \min_r \bar{\phi}(r)$ otherwise. In words, pairs of values such that $\beta_{1,j}\beta_{2,j} \geq 0$ are more heavily penalized by M_1, while all other pairs are equally penalized by Dirty and M_1. Consequently, assuming that $\beta^*_{1,j}\beta^*_{2,j} \geq 0$ for all $j \in [p]$, the probability of correct support recovery for M_1 is superior or equal to that of Dirty. In particular, under the assumptions considered in Theorem 3 of Jalali et al. (2013), and assuming that $\beta^*_{1,j}\beta^*_{2,j} \geq 0$ for all $j \in [p]$, perfect support recovery is ensured for method M_1 as soon as the common number of observations per stratum (task) is superior to $(2 - \alpha)s \log(p - (2 - \alpha)s)$, where s is the common support size of $\boldsymbol{\beta}^*_1$ and $\boldsymbol{\beta}^*_2$ and α is the overlap proportion between the two supports. As shown in Jalali et al. (2013) for their dirty model, this implies that M_1 outperforms both IndepLasso and the L_1/L_∞-group lasso strategy in terms of sample complexity (except in the two extremes situations of no sharing at all and full sharing, where it matches the best strategy).

6.1.2 Other preliminary non-asymptotic properties

In addition to the sample complexity in the simple case described above, other preliminary non-asymptotic properties can easily be derived for our approach thanks to its rewriting as a weighted lasso. Most non-asymptotic theoretical results of the lasso have been established under strong conditions on the design matrix: to name a few, the irrepresentability condition, the mutual incoherence, the restricted eigenvalue condition or the RIP property

(see van de Geer and Bühlmann (2009) for an overview). In Equation (5) of the main paper, it is clear that matrix $\boldsymbol{\mathcal{X}}$ generally does not enjoy such properties. Assuming for simplicity that $n_k = n/K$ for all $k \in [K]$, we present here simple situations where it does. A more thorough study would be needed to establish non-asymptotic properties under more general assumptions and an interesting lead lies in the recent results concerning prediction accuracy of lasso estimators under correlated designs (Dalalyan et al., 2014). This will be considered elsewhere.

The first situation is when K is large enough and the $\boldsymbol{\beta}_k^*$'s remain highly similar. More precisely, assume that design matrices $\mathbf{X}^{(k)}$ fulfill the following mutual incoherence condition (Lounici, 2008), for all $k \in [K]$: setting $\Sigma^{(k)} = ({\mathbf{X}^{(k)}}^T \mathbf{X}^{(k)})/n_k$ for all $k \in [K]$, we assume that $\Sigma^{(k)}_{j,j} = 1$ for all $j \in [p]$ and $\max_{j_1 \neq j_2} |\Sigma^{(k)}_{j_1,j_2}| \leq 1/(7\alpha s)$, for some integer $s \geq 1$ and constant $\alpha > 1$. Then the following standardized version of matrix $\boldsymbol{\mathcal{X}}$

$$\widetilde{\boldsymbol{\mathcal{X}}} = \begin{pmatrix} \mathbf{X}^{(1)} & \sqrt{K}\mathbf{X}^{(1)} & \dots & \mathbf{0} \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{X}^{(K)} & \mathbf{0} & \dots & \sqrt{K}\mathbf{X}^{(K)} \end{pmatrix}$$

can easily be shown to fulfill the same mutual incoherence assumption, as long as $K > (7\alpha s)^2$. For such values of K, and assuming that the true parameter vector $\boldsymbol{\theta}^*$ contains at most s non-zero components (such a $\boldsymbol{\theta}^*$ is ensured to be unique), results of Lounici (2008) can be applied to derive the L_∞ rate of convergence and the sign consistency of estimators obtained from M_1, with sparsity parameters set to $\lambda_{2,k} = \lambda_1$ (which is equivalent to working with the non-standardized matrix $\boldsymbol{\mathcal{X}}$ and $\lambda_{2,k} = \lambda_1 K^{-1/2}$). Of course the assumption $K > (7\alpha s)^2$ is very strong: it holds typically when strata are numerous and highly similar (so that s remains low).

Another simple situation arises when ${\mathbf{X}^{(k)}}^T \mathbf{X}^{(k)}/n_k = \mathbf{I}_p$ for all $k \in [K]$ and vectors $\boldsymbol{\beta}_k^*$ are all equal, and we set again $\lambda_{2,k} = \lambda_1$. Setting $J_0 = \{j : \theta_j \neq 0\}$, this ensures that $J_0 \subseteq [p]$, i.e. $\boldsymbol{\theta}^*_{J_0} = \overline{\boldsymbol{\beta}}^*_{J_0}$. Then it is easy to show that the irrepresentability condition of Wainwright (2009) holds with $\gamma = 1 - \sqrt{1/K}$, and then to obtain the following L_∞ bound: for any $\delta, \varepsilon > 0$, and assuming that the noise variables $\xi_i^{(k)}$ are i.i.d. σ-sub-Gaussian, we have

$$\max_k \|\widehat{\boldsymbol{\beta}}_k - \boldsymbol{\beta}_k^*\|_\infty = \|\widehat{\overline{\boldsymbol{\beta}}} - \overline{\boldsymbol{\beta}}^*\|_\infty \leq \sigma \left\{ \sqrt{\frac{2\log|J_0| + \varepsilon^2}{n}} + \frac{2}{\gamma}\sqrt{\frac{2|J_0|\log(Kp - |J_0|) + \delta^2}{n}} \right\},$$

with probability greater than $1 - 2\exp(-\delta^2/2) - 2\exp(-2\varepsilon^2/2)$. This bound, combined with a "beta-min" condition, leads to the sign consistency of the approach. It is noteworthy that the IdentLasso would lead to the same kind of bound, with an only slightly better second term: $2\sqrt{\{2|J_0|\log(p - |J_0|) + \delta^2\}/n}$. In other words, when ${\mathbf{X}^{(k)}}^T \mathbf{X}^{(k)}/n_k = \mathbf{I}_p$ for all $k \in [K]$ and all the $\boldsymbol{\beta}_k^*$'s are equal our approach performs as well as IdentLasso (up to constants and log-terms), which is optimal in this case.

6.1.3 Asymptotic results for the adaptive version of M_2 in the fixed p, fixed K case

In this paragraph, we will consider the simple situation where p and K are held fixed (they do not increase with n). We further assume that, for all $k \in [K]$, matrices ${\mathbf{X}^{(k)}}^T\mathbf{X}^{(k)}/n_k$ converges to positive definite matrices $\mathbf{C}^{(k)}$ as $n_k \to \infty$, and $n_k/n \to \kappa_k$ as $n \to \infty$, with $0 < \kappa_k < 1$ (i.e., *stratum* sizes all tend to infinity at the same rate). These assumptions essentially imply that MLEs $\widetilde{\beta}_{k,j}$ can be used in the definition of the weights in Equations (10) and (7) of the main paper, and are $\sqrt{n}$-consistent as $n \to \infty$.

Before stating our result, some notations are needed; in particular, the overall complexity of the model returned by M_2^{ad} has to be defined properly. For any $j \in [p]$, introduce the quantities $\ell_j^* = \min\{k : \beta_{k,j}^* \in \text{median}(\beta_{1,j}^*, \ldots, \beta_{K,j}^*)\}$, $K_j^* = \{k : \beta_{k,j}^* = \beta_{\ell_j^*,j}^*\} \subseteq [K]$, and $N_j^* = \sum_{k \in \ell_j^*} n_k$. Further introduce $\mathcal{A}^* = \{(k,j) : \beta_{k,j}^* \neq 0\}$ and, for all $j \in [p]$, $\mathcal{A}_j^* = \{k : \beta_{k,j}^* \neq 0\}$ and $K_{\mathcal{A}^*,j}^* = \{k : \beta_{k,j}^* = \beta_{\ell_j^*,j}^* \neq 0\}$.

To define the empirical counterparts of the previous quantities returned by M_2^{ad}, first recall that for any $j \in [p]$, $\ell_j = \min\{k : \widetilde{\beta}_{k,j} \in \text{median}(\widetilde{\beta}_{1,j}, \ldots, \widetilde{\beta}_{K,j})\}$. Then denote by $\widehat{\beta}_{k,j}^{(ad)}$ for any $k \in [K]$ and $j \in [p]$ the estimator of $\beta_{k,j}^*$ returned by M_2^{ad}. We can now define $\widehat{\mathcal{A}} = \{(k,j) : \widehat{\beta}_{k,j}^{(ad)} \neq 0\}$, and, for all $j \in [p]$, $\widehat{\ell}_j = \{k : \widehat{\beta}_{k,j}^{(ad)} = \widehat{\beta}_{\ell_j,j}^{(ad)}\} \subseteq [K]$, $\widehat{\mathcal{A}}_j = \{k : \widehat{\beta}_{k,j}^{(ad)} \neq 0\}$ and $\widehat{K}_{\widehat{\mathcal{A}},j} = \{k : \widehat{\beta}_{k,j}^{(ad)} = \widehat{\beta}_{\ell_j,j}^{(ad)} \neq 0\}$. Finally for any $j \in [p]$ denote by $\widehat{s}_j$ the number of distinct non-null values in $(\widehat{\beta}_{1,j}^{(ad)}, \ldots, \widehat{\beta}_{K,j}^{(ad)})$, that is

$$\widehat{s}_j = \begin{cases} |\widehat{\mathcal{A}}_j| & \text{if } \widehat{\beta}_{\ell_j,j}^{(ad)} = 0 \\ |\widehat{\mathcal{A}}_j| - |\widehat{K}_{\widehat{\mathcal{A}},j}| + 1 & \text{otherwise.} \end{cases}$$

The overall complexity of the model returned by M_2^{ad}, which corresponds to the number of "free" parameters returned by M_2^{ad}, is then defined as $\widehat{s} = \sum_{j \in [p]} \widehat{s}_j$. More precisely, for any $j \in [p]$ such that $\widehat{s}_j > 0$, M_2^{ad} returns a vector $\widehat{\boldsymbol{\eta}}_j^{(ad)}$ of size $\widehat{s}_j$ defined by

$$\widehat{\boldsymbol{\eta}}_j^{(ad)} = \begin{cases} \widehat{\beta}_{\ell_j,j}^{(ad)} & \text{if } \widehat{\beta}_{\ell_j,j}^{(ad)} \neq 0 \text{ and } \widehat{s}_j = 1; \\ (\widehat{\beta}_{\bar{k}_{j,1},j}^{(ad)}, \ldots, \widehat{\beta}_{\bar{k}_{j,\widehat{s}_j},j}^{(ad)})^T & \text{if } \widehat{\beta}_{\ell_j,j}^{(ad)} = 0 \text{ and } \widehat{s}_j \geq 1, \text{ with } \{\bar{k}_{j,1}, \ldots, \bar{k}_{j,\widehat{s}_j}\} = \widehat{\mathcal{A}}_j; \\ (\widehat{\beta}_{\ell_j,j}^{(ad)}, \widehat{\beta}_{\bar{k}_{j,1}^{(ad)},j}, \ldots, \widehat{\beta}_{\bar{k}_{j,\widehat{s}_j-1},j}^{(ad)})^T & \text{otherwise, with } \{\bar{k}_{j,1}, \ldots, \bar{k}_{j,\widehat{s}_j-1}\} = \widehat{\mathcal{A}}_j \setminus \widehat{K}_{\mathcal{A},j}. \end{cases}$$

Now denote by $\widehat{\boldsymbol{\eta}}^{(ad)}$ the estimator returned by M_2^{ad} (after some reordering), *i.e.*, $\widehat{\boldsymbol{\eta}}^{(ad)} = (\widehat{\boldsymbol{\eta}}_j^{(ad)})_{j:\widehat{s}_j>0} \in \mathbb{R}^s$. We can further introduce the theoretical counterpart of $\widehat{s}_j$ and $\widehat{s}$, by setting $s_j^* = |\mathcal{A}_j^*|$ if $\beta_{\ell_j^*,j}^* = 0$ and $s_j^* = |\mathcal{A}_j^*| - |K_{\mathcal{A}^*,j}^*| + 1$ otherwise, and $s^* = \sum_{j \in [p]} s_j^*$.

Finally introduce for any $j \in [p]$ such that $s^*_j > 0$,

$$\boldsymbol{\eta}^*_j = \begin{cases} \beta^*_{\ell^*_j,j} & \text{if } \beta^*_{\ell^*_j,j} \neq 0 \text{ and } s^*_j = 1; \\ (\beta^*_{\bar{k}^*_{j,1},j}, \ldots, \beta_{\bar{k}^*_{j,s_j},j})^T & \text{if } \beta^*_{\ell^*_j,j} = 0 \text{ and } s^*_j \geq 1, \text{ with } \{\bar{k}^*_{j,1}, \ldots, \bar{k}^*_{j,s_j}\} = \mathcal{A}^*_j; \\ (\beta^*_{\ell^*_j,j}, \beta^*_{\bar{k}^*_{j,1},j}, \ldots, \beta^*_{\bar{k}^*_{j,s^*_j-1},j})^T & \text{otherwise, with } \{\bar{k}^*_{j,1}, \ldots, \bar{k}^*_{j,s^*_j-1}\} = \mathcal{A}^*_j \setminus K^*_{\mathcal{A}^*,j}. \end{cases}$$

and $\boldsymbol{\eta}^* = (\boldsymbol{\eta}^*_j)_{j:s^*_j>0} \in \mathbb{R}^{s^*}$.

As a simple consequence of Theorem 2 in Viallon et al. (2014), we have $\mathbb{P}(\widehat{\mathcal{A}} = \mathcal{A}^*) \to 1$ and $\mathbb{P}(\cap_{j\in[p]}\{\widehat{K}_{\mathcal{A},j} = K^*_{\mathcal{A}^*,j}\}) \to 1$, as $n \to \infty$. In particular, this implies that the two vectors $\boldsymbol{\eta}^*$ and $\widehat{\boldsymbol{\eta}}^{(ad)}$ are of identical size $s^* = \widehat{s}$ with probability tending to one. Moreover, we have the following Gaussian limit distribution

$$\sqrt{n}(\widehat{\boldsymbol{\eta}}^{(ad)} - \boldsymbol{\eta}^*) \to \mathcal{N}(\mathbf{0}_{s^*}, \sigma^2(\widetilde{\mathcal{X}}^T_{s^*}\widetilde{\mathcal{X}}_{s^*})^{-1}) \quad as\ n \to \infty,$$

with $\widetilde{\mathcal{X}}_{s^*}$ the $(n \times s^*)$ matrix that can be written as $(\widetilde{\mathcal{X}}_{j,s^*_j})_{j:s^*_j>0}$ where each submatrix $\widetilde{\mathcal{X}}_{j,s^*_j}$, of size $(n \times s^*_j)$ is of the form:

$$\widetilde{\mathcal{X}}_{j,s^*_j} = \begin{cases} \sum_{k\in K^*_{\mathcal{A}^*,j}} \widetilde{X}^{(k)}_j & \text{if } \beta^*_{\ell^*_j,j} \neq 0 \text{ and } s^*_j = 1; \\ (\widetilde{X}^{(\bar{k}^*_{j,1})}_j, \ldots, \widetilde{X}^{(\bar{k}^*_{j,s^*_j})}_j) & \text{if } \beta^*_{\ell^*_j,j} = 0 \text{ and } s^*_j \geq 1, \text{ with } \{\bar{k}^*_{j,1}, \ldots, \bar{k}^*_{j,s_j}\} = \mathcal{A}^*_j; \\ (\sum_{k\in K^*_{\mathcal{A}^*,j}} \widetilde{X}^{(k)}_j, \widetilde{X}^{(\bar{k}^*_{j,1})}_j, \ldots, \widetilde{X}^{(\bar{k}^*_{j,s_j})}_j) & \text{otherwise, with } \{\bar{k}^*_{j,1}, \ldots, \bar{k}^*_{j,s^*_j-1}\} = \mathcal{A}^*_j \setminus K^*_{\mathcal{A}^*,j}. \end{cases}$$

Here, $\widetilde{X}^{(k)}_j$ is the column vector of size n of the form $(\mathbf{0}^T_{\sum_{k'<k} n_{k'}}, X^{(k)^T}_j, \mathbf{0}^T_{\sum_{k'>k} n_{k'}})^T$, with $X^{(k)}_j$ the j-th column of matrix $\mathbf{X}^{(k)}$. For any $j \in [p]$, if $\beta^*_{\ell^*_j,j} \neq 0$ then we have $\beta^*_{k,j} = \beta^*_{\ell^*_j,j}$ for all $k \in K^*_{\mathcal{A}^*,j}$ and these parameters are all estimated by the common value $\widehat{\beta}^{(ad)}_{\ell_j,j}$, which has the same Gaussian limit distribution as that of the estimator we would obtain by pooling all data corresponding to strata in $K^*_{\mathcal{A}^*,j}$ together. Our results then show that method M^{ad}_2 enjoys *asymptotic oracle properties*. More precisely, method M^{ad}_2 is asymptotically optimal in situations where, for all $j \in [p]$, $\beta^*_{k_1,j} = \beta^*_{k_2,j}$ implies that either $\beta^*_{k_1,j} = \beta^*_{k_2,j} = 0$ or $\beta^*_{k_1,j} = \beta^*_{k_2,j} = \beta^*_{\ell^*_j,j}$, as in examples 1 to 3 of Figure 4 of this Appendix. Asymptotic oracle properties can easily be derived for method M^{ad}_1 as well but, for M^{ad}_1 to be optimal we must have in addition: for all $j \in [p]$, $\{\beta^*_{\ell^*_j,j} \neq 0\} \Rightarrow \{\forall k \in [K], \beta^*_{k,j} \neq 0\}$. For instance, in example 2 on Figure 4, the estimator of $\beta^*_{1,j}$ returned by M^{ad}_1 can not be 0 (unless M^{ad}_1 returns a zero estimate for $\beta^*_{\ell_j,j}$ as well), so that M^{ad}_1 would typically return a model with overall complexity higher than the theoretical one. On the other hand, contrary

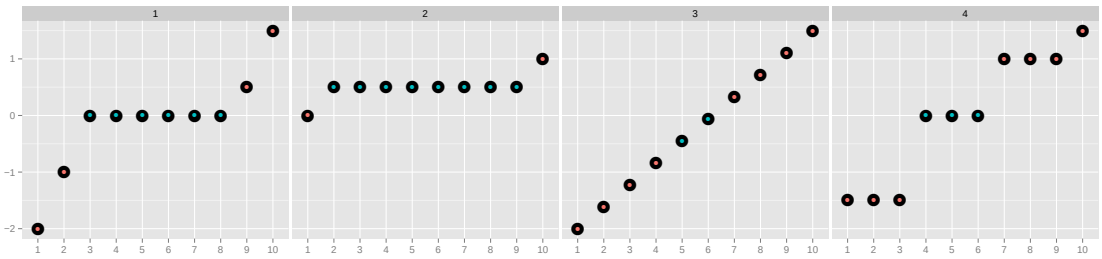

Figure 4: Graphical representation of four typical examples of vector $(\beta^*_{1,j}, \ldots, \beta^*_{K,j})$, for some $j \in [p]$ and with $K = 10$. Without loss of generality, we assume here that strata $\{1, \ldots, K\}$ are ordered in such a way that $\beta^*_{1,j} \leq \ldots \leq \beta^*_{K,j}$. For each example, the median value of $(\beta^*_{1,j}, \ldots, \beta^*_{K,j})$ is represented by a blue bullet: under the third example, all values between the two blue bullets $[-0.45, 0]$ are valid median values, while only one value corresponds to the median under the other three examples (0, 0.5 and 0 for examples 1, 2 and 4 respectively).

to the CliqueFused strategy that is based on cliques, both our approaches use star-graphs and are therefore sub-optimal in situations where non-zero values in $(\beta^*_{1,j}, \ldots, \beta^*_{K,j})$ consist of at least two groups of identical values (and possibly some distinct non-zero values). For instance, in example 4 of Figure 4, neither M_1^{ad} nor M_2^{ad} can return equal values for components $\widehat{\beta}^{(ad)}_{1,j}$, $\widehat{\beta}^{(ad)}_{2,j}$, and $\widehat{\beta}^{(ad)}_{3,j}$ (nor for $\widehat{\beta}^{(ad)}_{8,j}$, $\widehat{\beta}^{(ad)}_{9,j}$, and $\widehat{\beta}^{(ad)}_{10,j}$), while CliqueFused would, for n large enough. However, CliqueFused, *i.e.*, clique-based strategies, can of course be outperformed on finite samples. This was illustrated in the simulation study conducted in Viallon et al. (2014) in the single-stratum setting. In the present "multiple strata" setting, if $\beta^*_{1,j}, \ldots, \beta^*_{K-1,j}$ are all equal and $\beta^*_{K,j}$ is different from the $K-1$ other ones, we would have $\ell_j \in [K-1]$ for n enough, and our approaches would only penalize $|\beta_{K,j} - \beta_{\ell_j,j}|$: they are therefore more likely to detect this difference than the approach based on clique-graphs which penalizes all the differences $|\beta_{K,j} - \beta_{k,j}|$ for all $k \in [K-1]$.

6.2 Algorithm for method M_2

In this paragraph, we consider a general framework that encompasses both linear and logistic models along with SVM as special cases. Let f be a convex loss function and for any $k \in [K]$ set $\mathcal{L}_k(\boldsymbol{\beta}_k) = \sum_{i \in [n_k]} f(\boldsymbol{\beta}_k^T \mathbf{a}_i^{(k)} + c_i^{(k)})$ for $\boldsymbol{\beta}_k \in \mathbb{R}^p$ and some given $\mathbf{a}_i^{(k)} \in \mathbb{R}^p$ and $c_i^{(k)} \in \mathbb{R}$. We define the k-th *feature matrix* $\mathbf{A}^{(k)} := [\mathbf{a}_1^{(k)}, \ldots, \mathbf{a}_{n_k}^{(k)}]^T \in \mathbb{R}^{n_k \times p}$. We consider a generic supervised learning problem of the form

$$(\widehat{\overline{\boldsymbol{\beta}}}, (\widehat{\boldsymbol{\gamma}}_k)_{k=1}^K) \quad \in \quad \underset{\overline{\boldsymbol{\beta}}, (\boldsymbol{\gamma}_k)_{k=1}^K}{\operatorname{argmin}} \sum_{k=1}^K \Big\{ \mathcal{L}_k(\overline{\boldsymbol{\beta}} + \boldsymbol{\gamma}_k) + \lambda_{1,k} \|\overline{\boldsymbol{\beta}} + \boldsymbol{\gamma}_k\|_1 + \lambda_{2,k} \, \|\boldsymbol{\gamma}_k\|_1 \Big\}. \qquad (11)$$

	loss function $f(\xi)$	conjugate function $f^*(\vartheta)$	domain of f^*
squared	$f_{\text{sq}}(\xi) = \xi^2/(2n)$	$\vartheta^2/(2n)$	$\mathbb{R}$
logistic	$f_{\log}(\xi) = \log(1+e^{-\xi})$	$(-\vartheta)\log(-\vartheta) + (\vartheta+1)\log(\vartheta+1)$	$[-1,0]^m$
hinge	$f_{\text{hi}}(\xi) = (1-\xi)_+$	$-\vartheta$	$[-1,0]^m$

Table 1: Expression for the conjugate of popular loss functions, adopting the convention $0 \log 0 = 0$ for the logistic loss.

Recall that throughout the paper, data is of the form $(\mathbf{x}_1^{(k)}, y_1^{(k)}), \ldots, (\mathbf{x}_{n_k}^{(k)}, y_{n_k}^{(k)})$ with $\mathbf{x}_i^{(k)} \in \mathbb{R}^p$ and $y_i^{(k)} \in \mathbb{R}$, $i \in [n_k]$ and $k \in [K]$. In this context, the aforementioned formalism covers the linear regression with the loss function set to the squared loss: $f(\xi) = f_{\text{sq}}(\xi) = \xi^2/(2n)$, and by setting $\mathbf{a}_i^{(k)} = \mathbf{x}_i^{(k)} \in \mathbb{R}^p$ the vector of covariates and $c_i^{(k)} = -y_i^{(k)}$ the (negative) response for observation $i \in [n_k]$ of the k-th stratum. Likewise, by setting $f(\xi) = f_{\log}(\xi) = \log(1+e^{-\xi})$, $c_i^{(k)} = 0$, and $\mathbf{a}_i^{(k)} = y_i^{(k)} \mathbf{x}_i^{(k)}$ for $i \in [n_k]$, our formalism covers both the logistic regression and the SVM framework, by considering the logistic loss $f(\xi) = f_{\log}(\xi) = \log(1+e^{-\xi})$ and the hinge loss $f(\xi) = f_{\text{hi}}(\xi) = (1-\xi)_+$ respectively.

For future use, we denote by f^* the Fenchel conjugate of the loss function f, which is the extended-value convex function defined as

$$f^*(\vartheta) := \max_{\xi} \ \xi\vartheta - f(\xi).$$

Beyond convexity, we make a few mild assumptions about the loss function f. First, we assume that it is non-negative everywhere, and that it is closed (its epigraph is closed), so that $f^{**} = f$. These assumptions are met with the squared, logistic and hinge loss functions, as well as other popular loss functions. The conjugate of the squared, logistic and hinge loss functions are given in Table 1.

We have

$$\begin{aligned}
\phi((\lambda_{1,k})_{k\in[K]}, (\lambda_{2,k})_{k\in[K]}) &= \min_{\overline{\boldsymbol{\beta}}, (\boldsymbol{\gamma}_k)_{k=1}^K} \sum_{k=1}^K \big\{ \mathcal{L}_k(\overline{\boldsymbol{\beta}} + \boldsymbol{\gamma}_k) + \lambda_{1,k} \|\overline{\boldsymbol{\beta}} + \boldsymbol{\gamma}_k\|_1 + \lambda_{2,k} \|\boldsymbol{\gamma}_k\|_1 \big\} \\
&= \min_{\substack{\overline{\boldsymbol{\beta}}, (\boldsymbol{\beta}_k)_{k=1}^K, \\ (\mathbf{z}_k)_{k=1}^K}} \sum_{k=1}^K \big\{ G_k(\mathbf{z}_k) + \lambda_{1,k} \|\boldsymbol{\beta}_k\|_1 + \lambda_{2,k} \|\boldsymbol{\beta}_k - \overline{\boldsymbol{\beta}}\|_1 \big\} : \\
&\qquad \text{s.t.} \quad \mathbf{z}_k = \mathbf{A}^{(k)} \boldsymbol{\beta}_k + \mathbf{c}_k, k \in [K],
\end{aligned}$$

where, for $\mathbf{z} \in \mathbb{R}^{n_k}$, $G_k(\mathbf{z}) = \sum_{i=1}^{n_k} f(z_i)$ and $\mathbf{c}_k = (c_1^{(k)}, \ldots, c_{n_k}^{(k)})$ for any $k \in [K]$.

We can now express the problem in min-max form as follows

$$\phi((\lambda_{1,k})_{k\in[K]},(\lambda_{2,k})_{k\in[K]}) = \min_{\substack{\overline{\boldsymbol{\beta}},(\boldsymbol{\beta}_k)_{k=1}^K,\\ (\mathbf{z}_k)_{k=1}^K}} \max_{\substack{(\mathbf{u}_k)_{k=1}^K,(\mathbf{v}_k)_{k=1}^K,\\ (\boldsymbol{\alpha}_k)_{k=1}^K}} \sum_{k=1}^K \Big\{ G_k(\mathbf{z}_k) + \boldsymbol{\alpha}_k^T(\mathbf{A}^{(k)}\boldsymbol{\beta}_k + \mathbf{c}_k - \mathbf{z}_k) + \mathbf{u}_k^T\boldsymbol{\beta}_k + \mathbf{v}_k^T(\boldsymbol{\beta}_k - \overline{\boldsymbol{\beta}}) \Big\}$$
$$\text{s.t.} \quad \|\mathbf{u}_k\|_\infty \le \lambda_{1,k}, \|\mathbf{v}_k\|_\infty \le \lambda_{2,k}, k \in [K].$$

Now, solving for $\overline{\boldsymbol{\beta}}, (\boldsymbol{\beta}_k)_{k=1}^K$, we obtain the dual constraints:

$$\sum_{k=1}^K \mathbf{v}_k = \mathbf{0}_p, \quad {\mathbf{A}^{(k)}}^T \boldsymbol{\alpha}_k = \mathbf{u}_k + \mathbf{v}_k, \quad \text{for } k \in [K].$$

Still denoting by $f^\star$ the Fenchel conjugate of the function f, and eliminating variables u_k, $k = 1, \ldots, K$, we obtain the dual problem

$$\phi((\lambda_{1,k})_{k\in[K]},(\lambda_{2,k})_{k\in[K]})) = \min_{\boldsymbol{\alpha}_1,\ldots,\boldsymbol{\alpha}_K} \sum_k \Big\{ \boldsymbol{\alpha}_k^T\mathbf{c}_k + \sum_{i\in[n_k]} f^\star(\alpha_{k,i}) \Big\}$$

under the constraints:

$$\Big\| {\mathbf{A}^{(k)}}^T \boldsymbol{\alpha}_k - \mathbf{v}_k \Big\|_\infty \le \lambda_{1,k}, \ \|\mathbf{v}_k\|_\infty \le \lambda_{2,k}, \ \textstyle\sum_k \mathbf{v}_k = \mathbf{0}_p \text{ and } \alpha_{k,i} \in \mathcal{D}\mathrm{om}(f^\star), \text{ for } k \in [K].$$

For instance, in the logistic regression setting, this takes the form

$$\phi((\lambda_{1,k})_{k\in[K]},(\lambda_{2,k})_{k\in[K]})) = \min_{\boldsymbol{\alpha}_1,\ldots,\boldsymbol{\alpha}_K} \sum_{k\in[K]} \boldsymbol{\alpha}_k^T \log \boldsymbol{\alpha}_k + (1-\boldsymbol{\alpha}_k)^T \log(1-\boldsymbol{\alpha}_k)$$

under the constraints:

$$\Big\| {\mathbf{A}^{(k)}}^T \boldsymbol{\alpha}_k - \mathbf{v}_k \Big\|_\infty \le \lambda_{1,k}, \ \|\mathbf{v}_k\|_\infty \le \lambda_{2,k}, \ \textstyle\sum_k \mathbf{v}_k = \mathbf{0}_p, \text{ and } \boldsymbol{\alpha}_k \in [0,1]^{n_k}, \text{ for } k \in [K].$$

This problem is then equivalent to a standard entropy maximization and can therefore be solved using standard optimization toolboxes like the Mosek toolbox in R or Matlab for instance, which returns both optimal dual and primal solutions. (In the linear regression setting, the problem reduces to a quadratic programming and can also be solved with Mosek.)

6.3 (Sparse) group lasso (spGroupLasso).

Lounici et al. (2009) and Negahban and Wainwright (2011) consider two different group lasso strategies. Their estimators are defined as minimizers of the following criterion

$$\sum_k \frac{\|\mathbf{y}^{(k)} - \mathbf{X}^{(k)}\boldsymbol{\beta}^{(k)}\|_2^2}{2n} + \lambda\|\mathbf{B}\|_{1,q} \tag{12}$$

with typical values $q = 2$ (Lounici et al., 2009) or $q = \infty$ (Negahban and Wainwright, 2011). The $\|\cdot\|_{1,q}$ encourages solutions $\widehat{\mathbf{B}}$ to exhibit a row-wise group structure in the following sense: for each covariate $j \in [p]$, parameters $\widehat{\beta}_j^{(k)}$, for $k \in [K]$, are either all 0 or all non-zero (the estimated effect of each covariate is either null for all tasks or non-null for all tasks).

A slightly more flexible approach consists in adding an L_1 penalty to allow covariates to have a non-zero estimated effect on all but some tasks. This leads to the so-called sparse group lasso in standard regression models (Simon et al., 2013). In our context, this suggests to minimize to following criterion

$$\sum_k \frac{\|\mathbf{y}^{(k)} - \mathbf{X}^{(k)}\boldsymbol{\beta}^{(k)}\|_2^2}{2n} + \lambda_1\|\mathbf{B}\|_{1,1} + \lambda_2\|\mathbf{B}\|_{1,q}. \tag{13}$$

Note that in Equation (13) above solutions $\widehat{\mathbf{B}}$ are encouraged to be sparse and to exhibit a group structure while in the Dirty model of Jalali et al. (2013) (see Equation (6) of the main paper) solutions are sums of two matrices: $\widehat{\mathbf{S}}$ that is encouraged to be sparse and $\widehat{\mathbf{R}}$ that is encouraged to have a group structure.

Non-asymptotic properties have been established for the non-sparse group-lasso in multi-task settings (Lounici et al., 2009; Negahban and Wainwright, 2011). In particular, Negahban and Wainwright (2011) compare the L_1/L_∞-regularized method to IndepLasso in the special case of $K = 2$ linear regression problems with standard Gaussian designs whose supports have common size s and overlap in a fraction $\alpha \in [0, 1]$ of their entries. They focus on the capacity of each method to correctly specify the union of the two supports. Assuming a common number $n/2$ of observations for each task, they prove that L_1/L_∞-regularized method yields improved statistical efficiency if the overlap parameter is large enough ($\alpha \geq 2/3$), but has worse statistical efficiency than IndepLasso for moderate to small overlap ($\alpha < 2/3$). Recall that regarding sample complexity, the Dirty model was shown to strictly outperform both strategies under the same assumptions (except for extreme situations of full heterogeneity or full homogeneity where it matches the optimal strategy). Recall also that under the additional assumption $\beta^*_{1,j}\beta^*_{2,j} \geq 0$, we have established that our first approach M_1 performed at least similarly to the Dirty model (see Section 6.1.1 above).

6.4 Additional results from the simulation study

6.4.1 A simulation study when $K = 2$

Here we present a comparison of Dirty and M_1 when $K = 2$. We actually mimic the simulation study presented in Jalali et al. (2013). For a given value of p, the sparsity of both $\boldsymbol{\beta}_1^*$ and $\boldsymbol{\beta}_2^*$ is set to $s = \lfloor p/10 \rfloor$. Results are presented for $p = 128$ (results for $p = 256$ were similar). The proportion of non-zero and equal components in $\boldsymbol{\beta}_1^*$ and $\boldsymbol{\beta}_2^*$ is set to $\alpha \in \{0.5, 0.8\}$. For a given value of $\beta > 0$, non-zero components are set to $\pm\beta$ with probability $1/2$. Design matrices were generated under a multivariate Gaussian $\mathcal{N}_p(\mathbf{0}_p, \mathbf{I}_p)$ distribution, and noise variables under a $\mathcal{N}(0, 0.1)$ distribution. Number of observations per stratum $n_1 = n_2 = n/2$ were set to $n_k = cs \log[p-(2\alpha)s]/2$, for $c \in \{0.5, 1, 1.5, 2, 4\}$ defining the rescaled sample size (Jalali et al., 2013). Selection of the optimal sparsity parameter was performed by using an independent test sample. Methods referred to as 2step Dirty and 2step M_1 further use a 2-step strategy to unshrunk estimates before computing the L2 prediction error on the test sample. Figure 5 presents the results we obtained for the support recovery accuracy. It confirms that Dirty and M_1 perform similarly.

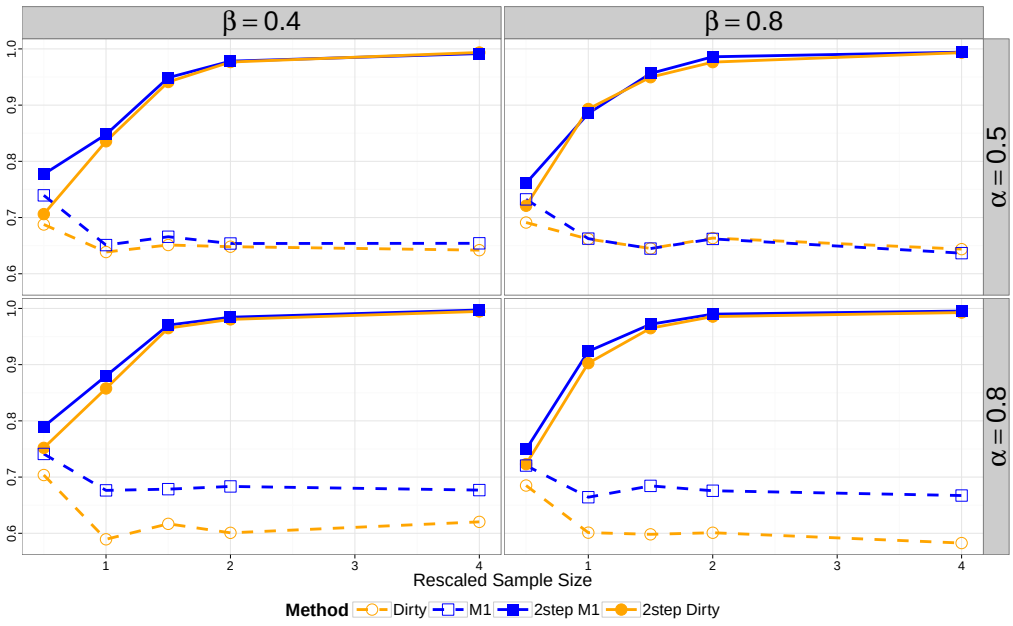

Figure 5: Results from the simulation study with $K = 2$.

6.4.2 Support recovery accuracy for the methods compared on Figure 1 of the main paper

Figure 6 presents the support recovery accuracies for the methods compared in terms of prediction accuracy on Figure 1 of the main paper. As mentioned in the main paper, conclusions drawn from Figure 6 are very similar to those drawn from Figure 1 of the main paper.

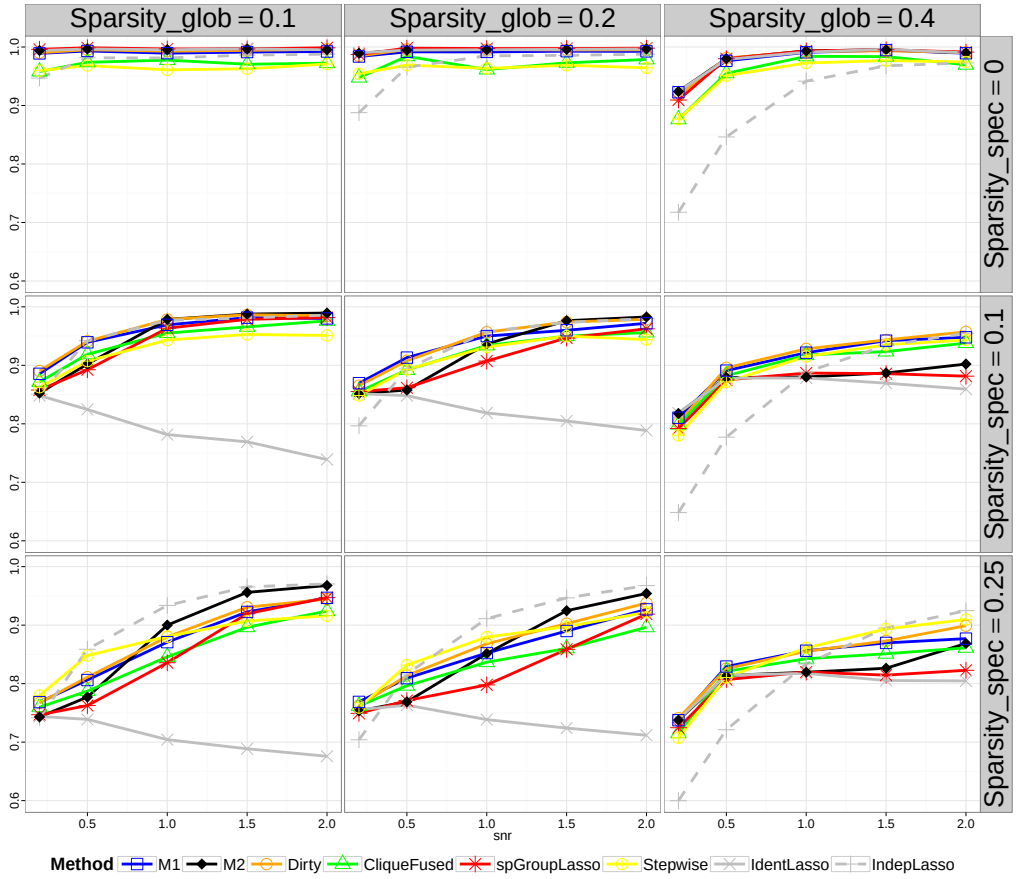

Figure 6: Results from the first simulation study: support recovery accuracy.

6.4.3 Performance of M_2 depending on the (λ_1, λ_2) grid

All methods presented in the main article involve tuning parameters that need to be carefully selected in practice. Generally speaking, a predefined grid of λ values has first to be constructed. For instance in the single stratum setting, the minimal value $\lambda_{\max}$ ensuring

that the vector returned by the lasso is null for all $\lambda \geq \lambda_{\max}$ has a closed-form expression (*e.g.*, for linear and logistic regression) and the grid is then generally of the form $\{\lambda_{\max}/1000, \ldots, \lambda_{\max}\}$. Most methods considered in the main paper involve two tuning parameters, and the derivation of $\lambda_{1,\max}$ and $\lambda_{2,\max}$ is not straightforward. Because our simulated data are such that $X_j^{(k)} \sim \mathcal{N}(0,1)$ for all $k \in [p]$ and $j \in [p]$, and $n_k = n/K$ for all $k \in [K]$, we limit the presentation to the case where tuning parameters do not depend on k; if the n_k's are not equal, one solution for M_1 consists in first standardizing the columns of matrix $\boldsymbol{\mathcal{X}}$ and then proceed as described here. For the Dirty approach of Jalali et al. (2013) we proceed as they did in their simulation study. Because of its connection with the lasso, the following strategy can be used for our first approach M_1 (or its adaptive version). For any pair (λ_1, λ_2) such that $\lambda_1 \geq K\lambda_2$, it is easy to see from Equation (4) of the main paper that the estimated common effect $\widehat{\overline{\boldsymbol{\beta}}}$ will be $\mathbf{0}_p$ and M_1 then reduces to a version of IndepLasso ran with the tuning parameter set to λ_2. On the other hand, if $\lambda_1 = 0$ then $\overline{\boldsymbol{\beta}}$ is unpenalized, and M_1 returns parameter vectors $\widehat{\boldsymbol{\beta}}_k = \widehat{\overline{\boldsymbol{\beta}}}$ for λ_2 large enough. Therefore, our strategy is to first make the λ_1/λ_2 ratio vary on the interval $[0, K]$ (we take 50 equally-spaced values on a log-scale), and for each value of this ratio r, the glmnet function can be used to compute the $\lambda_{2,\max}(r)$ value, along with the 50 models returned with $\lambda_2(r)$ varying on the grid of 50 equally-spaced values (on a log-scale) $\{\lambda_{2,\max}(r)/1000, \ldots, \lambda_{2,\max}(r)\}$.

For the other methods (M_2, CliqueFused and spGroupLasso) we proceed as follows. We first consider the case $\lambda_2 = 0$. The methods then all reduce to a standard lasso for which the $\lambda_{1,\max}$ value can be computed. The grid of λ_1-values is then chosen as the set of 50 values $\{\lambda_{1,\max}/1000, \ldots, \lambda_{1,\max}\}$, equally-spaced on a log-scale. Then, for each λ_1 value on this grid, the minimal value $\lambda_{2,\max}(\lambda_1)$ is numerically approximated. This value is such that for all $\lambda_2 \geq \lambda_{2,\max}(\lambda_1)$, the considered method computed with parameters (λ_1, λ_2) returns a vector of parameter $\widehat{\boldsymbol{\beta}}(\lambda_1, \lambda_2)$ with the lowest possible overall complexity. In particular, for methods M_1, M_2 and CliqueFused, this means that $\widehat{\boldsymbol{\beta}}_k(\lambda_1, \lambda_2) = \widehat{\boldsymbol{\beta}}_{k'}(\lambda_1, \lambda_2)$ for all $(k, k') \in [K]$ and $\lambda_2 \geq \lambda_{2,\max}(\lambda_1)$. Such an approximate $\lambda_{2,\max}(\lambda_1)$ value is simply obtained by iteratively running the considered method for various values of λ_2, starting from a huge value, and successively dividing it by 2 for instance. Another, more simple, strategy is as follows. We first consider the case $\lambda_2 = 0$. The methods then all reduce to a standard lasso for which the $\lambda_{1,\max}$ value can be computed. Then, we use the grid of 50 values $\{\lambda_{1,\max}/1000, \ldots, \lambda_{1,\max}\}$ for both λ_1 and λ_2. This strategy is referred to as AlternativeGrid on Figure 7 below which presents the comparison of the performance for M_2 implemented with the two strategies on the first simulation study. The comparison of the black and grey curves illustrates how difficult it is to empirically compare methods. When using the AlternativeGrid strategy, M_2 performs similarly to M_1, while using the other strategy, it performs like SpGroupLasso (which uses the same strategy for the grid construction). Therefore, the observed difference of performance between the various

strategies can be (at least partly) explained by the choice for the initial grid of (λ_1, λ_2) values.

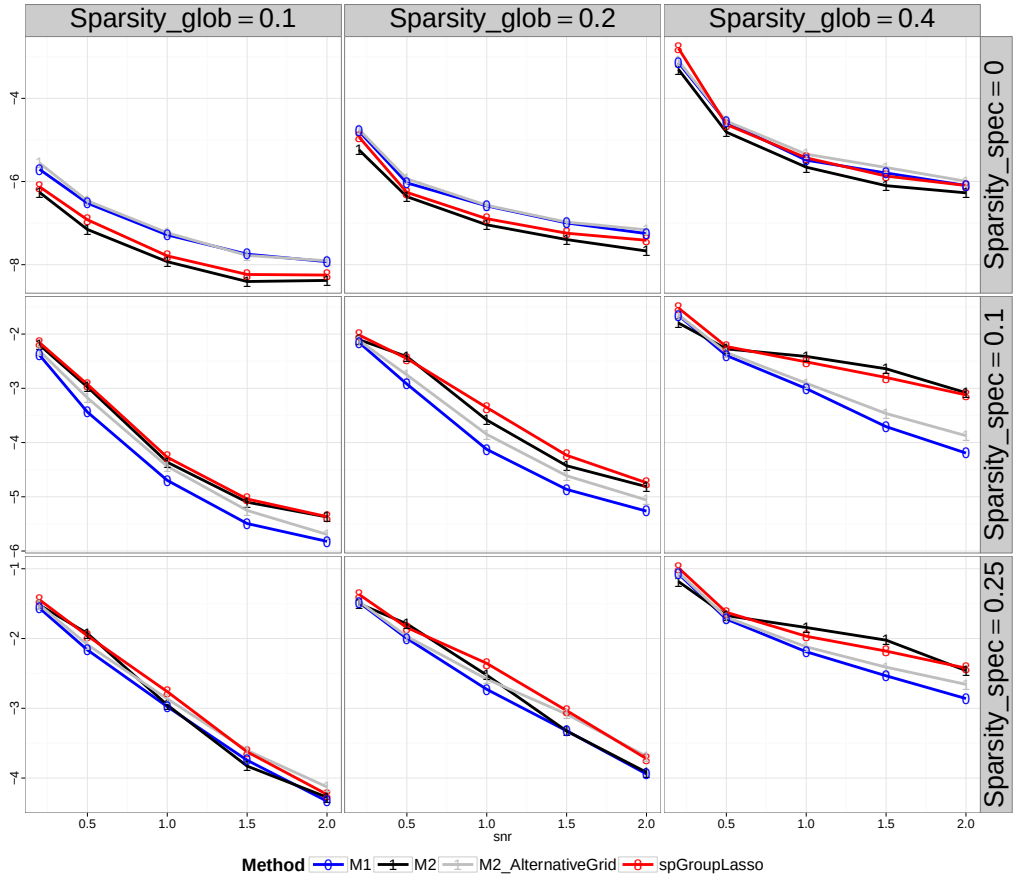

Figure 7: Comparison of the prediction performance for M_2 depending on the strategy used to construct the (λ_1, λ_2) grid of values.